Flexible online multivariate regression with variational Bayes and the matrix-variate Dirichlet process

Meng Hwee Victor Ong, David J Nott and Ajay Jasra* †

arXiv:1602.08849v1 [stat.CO] 29 Feb 2016

SUMMARY

Flexible regression methods where interest centres on the way that the whole distribution of a response vector changes with covariates are very useful in some applications. A recently developed technique in this regard uses the matrix-variate Dirichlet process as a prior for a mixing distribution on a coefficient in a multivariate linear regression model. The method is attractive, particularly in the multivariate setting, for the convenient way that it allows for borrowing strength across different component regressions and for its computational simplicity and tractability. The purpose of the present article is to develop fast online variational Bayes approaches to fitting this model and to investigate how they perform compared to MCMC and batch variational methods in a number of scenarios.

Keywords: Bayesian nonparametrics; Dirichlet process; Matrix-variate Dirichlet process; Variational Bayes.

1 Introduction

Flexible modelling of multivariate conditional densities is a fundamental problem in statistics, particularly in regression applications in which there is interest in the ways that the whole distribution of a response vector depends on covariates. In a recent paper Zhang et al. (2010) developed a flexible multivariate regression method using a Dirichlet process prior for a mixing distribution on the coefficient in a multivariate linear model, where the Dirichlet process base prior is a matrix-variate normal distribution. The approach is attractive for its flexibility, the easy way it allows borrowing of strength between regressions for different response variables through the matrix-variate normal base prior, and the computational simplicity and convenience that comes from basing the method

* Victor Ong is Research Associate, Department of Statistics and Applied Probability, National University of Singapore, Singapore 117546. (email : victor84@nus.edu.sg). David J. Nott is Associate Professor, Department of Statistics and Applied Probability, National University of Singapore, Singapore 117546. (email :standj@nus.edu.sg). Ajay Jasra is Associate Professor, Department of Statistics and Applied Probability, National University of Singapore, Singapore 117546 (email : staja@nus.edu.sg).

†Victor Ong, David J. Nott and Ajay Jasra were supported by a Singapore Ministry of Education Academic Research Fund Tier 2 grant (R-155-000-143-112).

on the ordinary Dirichlet process. They refer to the Dirichlet process prior with matrix-variate normal base measure as the matrix-variate Dirichlet process (hereafter MDP), and further applications beyond the multivariate linear regression setup were considered in Zhang et al. (2014). The contribution of the present work is to consider fast online approaches to fitting the model of Zhang et al. (2010) using variational Bayes methods, suitable for application in the context of large datasets. We also consider a novel approach to improving the predictive performance of the online algorithm which gives performance comparable in many cases to a batch variational Bayes or MCMC approach.

In Bayesian nonparametrics, the development of suitable prior distributions for regression problems of the kind we consider here, involves the development of dependent prior distributions for sets of distributions indexed by the covariates. A recent survey on the extensive literature on this topic is given by Foti and Williamson (2015). A key early paper is by MacEachern (2000), who introduced the framework of the dependent Dirichlet process and which inspired many later developments. Some of the existing approaches in the literature include starting from the stick breaking representation of a random measure and allowing distribution atoms or weights to be covariate dependent (De Iorio et al., 2004; Gelfand et al., 2005; Griffin and Steel, 2006; Dunson and Park, 2008); consideration of covariate dependent generalizations of the Chinese restaurant process or Pólya urn prediction rule (Blei and Frazier, 2011; Caron et al., 2007); as well as methods that build on normalized completely random measures (Kingman, 1967; Lijoi and Prünster, 2010) and which use their relationship with Poisson processes to introduce covariate dependence in various ways (Rao and Teh, 2009; Chen et al., 2013; Lijoi et al., 2014). The above list of references is by no means exhaustive. For the special case of grouped data, the hierarchical Dirichlet process (Teh et al., 2006) has also been an extremely important development.

As mentioned, in the present work we consider the model of Zhang et al. (2010) which is attractive in the case of multivariate response for the convenient mechanism it represents for borrowing strength across regressions for different components through the matrix-variate normal base prior. Our objective is to develop fast online variational Bayes methods which allow the model of Zhang et al. (2010) to be applied with large datasets. The approach adopted builds on the VSUGS algorithm of Zhang et al. (2014) for Dirichlet process mixture models, which is a variational extension of the SUGS algorithm of Wang and Dunson (2011). Lin (2013) independently developed a similar algorithm to that of Zhang et al. (2014). The development of fast variational methods for complex Bayesian nonparametric models has been a very active area of recent research, with an important early paper being Blei and Jordan (2006) where a batch variational algorithm for fitting Dirichlet process mixture models was developed. In the online setting, some

recent contributions include Wang et al. (2011) and Bryant and Sudderth (2012) who consider online algorithms for the hierarchical Dirichlet process, and various methods inspired by the stochastic variational inference framework of Hoffman et al. (2013) (for example, Wang and Blei (2012)). Kabisa et al. (2016) consider a fast online approach to fitting high-dimensional correlated data with a model incorporating some Bayesian nonparametric components; their method is a variational Bayes algorithm which is similar in approach to methods originally developed by Sato (2001). Luts et al. (2014) consider online approaches to fitting semiparametric regression models in the variational Bayes framework.

The next section describes the matrix-variate Dirichlet process mixture model that is considered throughout the rest of the article. In Section 3, a batch variational algorithm for the model is derived and then Section 4 discusses the VSUGS online algorithm which is able to work efficiently for very large datasets. Section 5 discusses predictive inference and our novel regression adjustment approach. Section 6 considers an application to weak informative prior selection, Section 7 considers predictive performance of the methods in some benchmark data sets and Section 8 concludes.

2 Matrix-variate Dirichlet process mixture model

We consider the matrix-variate Dirichlet process mixture model of Zhang et al. (2010). Specifically, let y_i, $i = 1, \ldots, n$ denote a collection of observed m-dimensional response vectors and x_i, $i = 1, \ldots, n$ denote corresponding p-dimensional vectors of covariates. A common flexible way to model the mean in a multivariate regression for the responses involves using some basis expansion where, denoting the jth element of y_i by y_{ij},

$$E(y_{ij}) \quad = \quad \beta_{0,j} + \sum_{r=1}^{N} \beta_{r,j} E_r(x_i) \tag{1}$$

where $E_r(x)$, $r = 1, \ldots, N$ are basis functions and $\beta_j = (\beta_{0j}, \ldots, \beta_{Nj})^T$ are coefficients, $j = 1, \ldots, m$. In motivating their approach Zhang et al. (2010) discuss such a basis expansion, and consider setting $N = n$ and $E_r(x) = K(x, x_r)$ where $K(\cdot, \cdot)$ is a kernel function so that the number of basis terms equals the number of observations. Here we will be concerned with an online implementation of their approach where n is not known beforehand, so we will make a fixed choice of both N and the basis functions $E_r(x)$, $r = 1, \ldots, N$. We give more details about this later.

Write $\beta = [\beta_1, \ldots, \beta_m]$ for the $(N + 1) \times m$ matrix of regression coefficients and $E_i = (1, E_1(x_i), \ldots, E_N(x_i))^T$. Then if we assume i.i.d errors in the regression (1) we can

write

$$y_i = \beta^T E_i + \epsilon_i$$

where the ϵ_i are the errors having mean 0 and covariance matrix $\tau\Sigma$ say where $\tau > 0$ is a scale parameter. The reason for parametrizing the covariance matrix in this way will become clear later when conjugate prior specifications are considered. Flexible multivariate regression approaches using basis expansions of this type have been considered by many authors. The innovation of Zhang et al. (2010) is to consider a model in which the coefficient β varies randomly between observations. The distribution of this coefficient is estimated from the data, and is given a Dirichlet process prior with a matrix-variate normal distribution as the base measure. That is, the Dirichlet process with matrix-variate normal base measure is used as a prior on the mixing distribution for the coefficient. The clustering property of the Dirichlet process ensures that many observations will share the same coefficient matrix and there is borrowing of strength both between observations and responses in estimating the regression.

Precisely, the model is

$$\begin{aligned} Y_i|E_i, \tilde{\beta}_i, \Sigma &\sim N(\tilde{\beta}_i^T E_i, \tau\Sigma) \\ \tilde{\beta}_i|Q &\sim Q \\ Q|\alpha, M &\sim DP(\alpha, M) \end{aligned} \tag{2}$$

where $DP(\alpha, M)$ denotes the Dirichlet process with precision parameter α and base measure M. The base measure M in the model is chosen to be a matrix-variate normal distribution $N_{N+1,m}(0, \Omega \otimes \Sigma)$. An $s \times t$ random matrix Z has a matrix-variate normal distribution $N_{s,t}(C, V \otimes W)$, where C is an $s \times t$ matrix and V and W are $s \times s$ and $t \times t$ covariance matrices respectively, if its density takes the form

$$p(Z) = (2\pi)^{-st/2}|V|^{-t/2}|W|^{-s/2} \exp\left(\text{tr}\left(-\frac{1}{2}V^{-1}(Z-C)W^{-1}(Z-C)^T\right)\right).$$

In our model following Zhang et al. (2010) it will be assumed that Ω is diagonal, $\Omega = \text{diag}(\omega_1, \ldots, \omega_{N+1})$ where $\omega_i \sim IG(a_i, b_i)$ with a_i and b_i known. Also, Σ is inverse-Wishart with degrees of freedom ν and scale matrix S. τ is given an inverse gamma prior $IG(a_\tau, b_\tau)$ with a_τ and b_τ known.

The Dirichlet process puts all its mass on a countable collection of points so we can rewrite the model in the following way. Let $\{\beta_i\}_{i=1}^{\infty}$ be the distinct values appearing in the sequence $\{\tilde{\beta}_i\}_{i=1}^{\infty}$ with the β_i indexed according to their order of occurrence in $\{\tilde{\beta}_i\}_{i=1}^{\infty}$. We let δ_i be an integer valued variable with $\delta_i = j$ if $\tilde{\beta}_i = \beta_j$. Write $\delta_{1:i} = (\delta_1, \ldots, \delta_i)^T$. Using the Pólya urn representation for the Dirichlet process we can rewrite the model in

the form

$$\begin{aligned} Y_i|E_i, \beta, \Sigma, \delta_i &\sim N(\beta_{\delta_i}^T E_i, \tau\Sigma) \\ p(\delta, \beta) &= p(\delta)p(\beta) \end{aligned} \tag{3}$$

where $p(\delta_{1:n}) = p(\delta_{1:n}|\alpha) = \prod_{i=1}^{n} p(\delta_i|\delta_{1:i-1}, \alpha)$, $p(\beta) = \prod_{i=1}^{\infty} p(\beta_i)$ with $p(\beta_i)$ the matrix-variate normal density $N_{N+1,m}(0, \Omega \otimes \Sigma)$, the priors on Ω and Σ are the same as before and the conditional densities $p(\delta_i|\delta_{1:i-1}, \alpha)$ are defined by (using similar notation to Zhang et al. (2010))

$$p(\delta_i = j|\delta_{1:i-1}, \alpha) = \begin{cases} \frac{n_j^{(i)}}{\alpha+i-1} & j \in \{1, \ldots, n_i\} \\ \frac{\alpha}{\alpha+i-1} & j = n_i + 1 \end{cases}$$

where $n_j^{(i)}$ is the number of δ_k, $k < i$ equal to j and n_i is the number of distinct β_k appearing up to time $i - 1$. For the purpose of developing our fast online variational approximation algorithm we will use a truncated Dirichlet process mixture model. In this model the sequence $\{\beta_i\}_{i=1}^{\infty}$ is truncated to $\{\beta_i\}_{i=1}^{T}$ where T is the truncation point and

$$p(\delta_i = j|\delta_{1:i-1}, \alpha) = \begin{cases} \frac{n_j^{(i)}+\alpha/T}{\alpha+i-1} & j \in \{1, \ldots, n_i\} \\ \frac{\alpha(1-n_i/T)}{\alpha+i-1} & j = n_i + 1 \end{cases} . \tag{4}$$

This is the model we discuss in what follows.

3 Variational inference

Consider a Bayesian model with parameter ξ, prior $p(\xi)$ and likelihood $p(y|\xi)$. Variational Bayes computational methods (Waterhouse et al., 1996; Jordan et al., 1999; Attias, 2000; Ormerod and Wand, 2010) attempt to approximate the posterior density $p(\xi|y)$ by a more tractable and manageable variational density $q(\xi)$, belonging to a convenient family. The choice of $q(\xi)$ within the approximating family is usually made by minimizing the KL divergence between $p(\xi|y)$ and $q(\xi)$. It can be shown that

$$\log p(y) = \int \log\left(\frac{p(\xi)p(y|\xi)}{q(\xi)}\right) q(\xi)d\xi + \int \log\left(\frac{q(\xi)}{p(\xi|y)}\right) q(\xi)d\xi \tag{5}$$

where $p(y) = \int p(\xi)p(y|\xi)\, d\xi$. The first and second terms on the RHS of (5) are the variational lower bound $\mathcal{L}$ (so-called because it forms a lower bound on $\log p(y)$) and the KL divergence between $q(\xi)$ and $q(\xi|y)$, respectively. From (5), it is clear that minimizing the KL divergence is equivalent to maximizing $\mathcal{L}$. For further background see the references above.

Now, suppose that ξ can be partitioned into J subvectors, $\xi_1, .., \xi_J$. In variational Bayes, an approximating family for the posterior is considered where $q(\xi)$ is assumed to factorize as $\prod_{j=1}^{J} q(\xi_j)$. For each of the factors $q(\xi_j)$, the lower bound is maximized with the other factors held fixed by choosing $q(\xi_j)$ as

$$\hat{q}(\xi_j) \propto \exp\left\{E_{-\xi_j} \log p(y|\xi)p(\xi)\right\} \tag{6}$$

where $E_{-\xi_j}$ denotes an expectation with respect to $\prod_{i \neq j} q(\xi_i)$. Expression (6) is the basis of a blockwise gradient descent algorithm for maximizing $\mathcal{L}$ where an initial choice is made for the factors and then each factor is updated in turn with the others fixed at current values until convergence.

One useful application of the variational approach is to approximate the posterior distribution of parameters in Bayesian nonparametric models. It is well known that there is usually no direct way to compute the posterior distribution in these models and that MCMC sampling methods for such models can be difficult and computationally expensive. These considerations motivated Blei and Jordan (2006) to consider a mean-field variational inference algorithm for Dirichlet process mixture models. Their approach can be implemented for the model of Section 2, since the approach of Zhang et al. (2010) is based on an ordinary Dirichlet process mixture model, and we do implement such an approach later in our examples. Since this is a straightforward application of the algorithm of Blei and Jordan (2006) we do not give further details of their method here. However, we develop an alternative batch variational Bayes algorithm which is also described in the next section. The algorithm of Blei and Jordan (2006) is based on the stick breaking representation of the Dirichlet process; our alternative batch variational Bayes algorithm (like the later sequential algorithm of Section 4) is based on the Pólya urn representation with the unknown mixing distribution integrated out. Although the alternative batch algorithm involves some further approximations, the purpose of developing this method is that it gives a batch algorithm similar to our later online approach, and provides another reference for comparison for the performance of the online algorithm where how much performance is lost through the sequential updating mechanism can be better understood. Also, many of the updating steps in the online algorithm are simple modifications of the corresponding steps for the batch algorithm.

3.1 Batch mean field updates for global parameters

We work with the model (3). For the matrix-variate DP mixture model, variational inferences for the parameters $\beta_{1:T}, \Sigma, \tau, \omega_{1:N}, \delta_{1:n}$ are required. Define $\theta = (\beta_{1:T}, \Sigma, \tau, \omega_{1:N})^T$ and $\theta_j = (\beta_j, \Sigma, \tau, \omega_{1:N})^T$ for $j = 1, .., T$. In deriving approximate mean field updates we

consider a slight expansion of the model (3). This will be helpful when discussing the online case later, since in our model expansion the variational posterior has the same form as the prior leading to a natural online implementation. In the expanded model the prior on β_i is changed from $N_{N+1,m}(0, \Omega \otimes \Sigma)$ to $N_{N+1,m}(M_i, \Omega_i \otimes \Sigma)$ with $\Omega_i = (\Omega^{-1} + C_i)^{-1}$ where M_i and C_i are known matrices. Letting $\beta = (\beta_1, \ldots, \beta_T)$, we consider the following factorization for the variational posterior distribution:

$$q(\theta, \delta_{1:n}) = q(\beta, \Sigma)q(\tau)q(\omega_{1:N+1})q(\delta_{1:n}).$$

$q(\delta_i = j)$ will be denoted by q_{ij}. In this subsection we give the mean field updates for all factors except for $q(\delta_{1:n})$, which is considered in the next subsection. Technical details of the derivations are found in Appendix A.

For β, we recognize the form of $q(\beta, \Sigma)$ as being $q(\beta, \Sigma) = q(\Sigma)q(\beta|\Sigma)$ where $q(\Sigma)$ is inverse Wishart, and $q(\beta|\Sigma) = \prod_{j=1}^{T} q(\beta_j|\Sigma)$ with $q(\beta_j|\Sigma) = N_{N+1,m}(\hat{\beta}_j, V_j^{-1} \otimes \Sigma)$,

$$\hat{\beta}_j = V_j^{-1}\left((E_q(\Omega^{-1}) + C_j)M_j + E_q(\tau^{-1})\sum_{i=1}^{n} q_{ij}E_i y_i^T\right)$$

and

$$V_j = \left(E_q(\Omega^{-1}) + C_j + E_q(\tau^{-1})\sum_{i=1}^{n} q_{ij}E_i E_i^T\right).$$

For Σ, $q(\Sigma) = IW(\hat{\nu}, \hat{S})$ where

$$\hat{\nu} = \nu + n, \hat{S} = S + \sum_{j=1}^{T}\left(E_q(\tau^{-1})\sum_{i=1}^{n} q_{ij}y_i y_i^T + M_j^T(E_q(\Omega^{-1}) + C_j)M_j - \hat{\beta}_j^T V_j \hat{\beta}_j\right).$$

For τ, $q(\tau) = IG(\hat{a}_\tau, \hat{b}_\tau)$ where

$$\hat{a}_\tau = a_\tau + \frac{nm}{2}, \hat{b}_\tau = b_\tau + \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{T} q_{ij}((y_i - \hat{\beta}_j^T E_i)^T E_q(\Sigma^{-1})(y_i - \hat{\beta}_j^T E_i) + mE_i^T V_j^{-1} E_i).$$

Lastly, for $\omega = (\omega_1, .., \omega_{N+1})$, $q(\omega) = \prod_{i=1}^{N+1} q(\omega_i)$. Each $q(\omega_i) = IG(\hat{a}_i, \hat{b}_i)$ where

$$\hat{a}_i = a_i + mT/2, \hat{b}_i = b_i + \frac{1}{2}\sum_{j=1}^{T}\left((\hat{\beta}_{j,i} - M_{j,i})E_q(\Sigma^{-1})(\hat{\beta}_{j,i} - M_{j,i})^T + m\omega'_{ij}\right). \quad (7)$$

where $\hat{\beta}_{j,i}$ is the ith row of $\hat{\beta}_j$ and ω'_{ij} is the ith diagonal element of V_j^{-1}.

3.2 Batch mean field update for local parameters

We now factorize $q(\delta_{1:n})$ as $\prod_{i=1}^{n} q(\delta_i)$ and consider approximate mean field updates for $q(\delta_i)$, $i = 1, \ldots, n$. Using (6), for each δ_i, we get

$$q(\delta_i) \propto \exp\left(E_q\left\{\log p(\delta_i|\delta_{\neq i}, \alpha)\right\} + E_q\left\{\log(y_i|\theta, \delta_i)\right\}\right).$$

where $\delta_{\neq i}$ denotes $\delta_{1:n}$ with δ_i omitted. Making the approximation

$$\exp\left(E_q\left\{\log p(\delta_i|\delta_{\neq i}, \alpha)\right\}\right) \approx E_q\left\{p(\delta_i|\delta_{\neq i}, \alpha)\right\}$$

we have

$$q(\delta_i) \propto E_q\left\{p(\delta_i|\delta_{\neq i}, \alpha)\right\} \exp\left(E_q\left\{\log(y_i|\theta, \delta_i)\right\}\right).$$

If $i \geq T$, we approximate further $E_q\left\{p(\delta_i = j|\delta_{\neq i}, \alpha)\right\}$ by

$$\frac{\sum_{k \neq i} q_{kj} + \alpha/T}{\alpha + n - 1},$$

with the case where $i < T$ being handled by using the same expression but conditioning on $\delta_i \leq i$. This approximation is obtained by reordering so that the ith observation is last, taking an expectation in (4) and then restoring the constraint associated with the original ordering by conditioning on $\delta_i \leq i$ if $i \leq T$. Note that because we order atoms according to their order of occurrence it must be the case that $\delta_i \leq i$ for $i \leq T$. To get an expression for our approximate mean field update it remains to evaluate $E_q\left\{\log(y_i|\theta_j)\right\}$ which is

$$\begin{aligned}
E_q\left\{\log(y_i|\theta_j)\right\} &= -\frac{m}{2}\log(2\pi) - \frac{m}{2}E_q\left\{\log\tau\right\} - \frac{1}{2}E_q\left\{\log|\Sigma|\right\} \\
&\quad - \frac{1}{2}E_q\left\{\frac{1}{\tau}\right\} E_q\left\{y_i^T\Sigma^{-1}y_i - 2E_i^T\beta_j\Sigma^{-1}y_i + E_i^T\beta_j\Sigma^{-1}\beta_j^T E_i\right\} \\
&= -\frac{m}{2}\log(2\pi) - \frac{m}{2}\left\{\log(\hat{b}_\tau) - \psi(\hat{a}_\tau)\right\} \\
&\quad - \frac{1}{2}\left\{-\psi_m\left(\frac{\hat{\nu}}{2}\right) - m\log(2) + \log|\hat{S}|\right\} - \frac{1}{2}\frac{\hat{a}_\tau}{\hat{b}_\tau}\left\{\hat{\nu}y_i^T(\hat{S})^{-1}y_i\right. \\
&\quad \left. -2\hat{\nu}E_i^T\hat{\beta}_j(\hat{S})^{-1}y_i + mE_i^T\hat{\Omega}_j E_i + (E_i^T\hat{\beta}_j)^T(\hat{\nu}(\hat{S})^{-1})(E_i^T\hat{\beta}_j)\right\}.
\end{aligned}$$

4 VSUGS for matrix-variate Dirichlet process mixture model

The VSUGS algorithm, proposed by Zhang et al. (2014), is an online learning procedure for fast fitting of Dirichlet process mixture models. It uses the variational approximation framework to improve the SUGS algorithm (Wang and Dunson, 2011). The VSUGS algorithm is especially useful for large datasets as computing the full variational batch update or using MCMC might be computationally infeasible. The framework of the VSUGS procedure is as follows. Following Zhang et al. (2014), we consider an approximation to the posterior $p(\delta_{1:i-1}, \theta_{1:T}|y_{i:i-1})$ of the form

$$\prod_{j=1}^{i-1} q_{i-1}(\delta_j) \prod_{j=1}^{T} q_{i-1}(\theta_j).$$

The algorithm starts at $\hat{q}_1(\delta_1 = 1) = 1$, $\hat{q}(\theta_1) = p(\theta_1|y_1, \delta_1 = 1)$. Then, at time i, we use $\hat{q}_{i-1}(\theta)$ and $\hat{q}_{i-1}(\delta_{1:i-1})$ as a prior for processing the data point y_i. Then for a certain fixed choice of $\hat{q}_i(\delta_i)$ the mean field update for θ reduces to the following approximation of $p(\theta|y)$:

$$\hat{q}_i(\theta) \propto \hat{q}_{i-1}(\theta) \prod_{j=1}^{T} \exp\left(\hat{q}_i(\delta_i = j) E_q\left\{\log p(y_i|\theta_j)\right\}\right). \tag{8}$$

For the assignment variables δ_i, we follow Zhang et al. (2014) and choose

$$\hat{q}_i(\delta_i = j) = r_{ij} \int \hat{q}_{i-1}(\theta_{\delta_i}) p(y_i|\theta_{\delta_i}) d\theta_{\delta_i} \tag{9}$$

for $j \in \{1, ..., \min(i, T)\}$ where T is the pre-specified truncation point for the number of mixture components and

$$r_{ij} = \begin{cases} \frac{\sum_{k=1}^{i-1} \hat{q}_{i-1}(\delta_k = j) + \alpha/T}{\alpha + i - 1} & j \in \{1, \ldots, \min(i-1, T)\}, \\ \frac{\alpha(1 - \min(i-1, T)/T)}{\alpha + i - 1} & j = \min(i-1, T) + 1. \end{cases} \tag{10}$$

One property of the VSUGS procedure is that (9) splits the likelihood contribution from the ith observation among the mixture components. This deviates from the original SUGS algorithm (Wang and Dunson, 2011) which uses a "hard" allocation to mixture components. In the case of conjugate priors, the VSUGS algorithm retains the computational advantages of the original SUGS algorithm. See Zhang et al. (2014) for further details.

4.1 Sequential update of variational parameters for τ, Σ and Ω

In the batch update of the global parameters the expectations of τ^{-1}, Σ^{-1} and Ω^{-1} with respect to $q(\tau)$, $q(\Sigma)$ and $q(\Omega)$ respectively are required. For an online algorithm like VSUGS, these expectations change when a new data point enters. In order to use (8), it is required to replace the expectation of τ^{-1}, Σ^{-1} and Ω^{-1} with $E_{q_i}(\tau^{-1}), E_{q_i}(\Sigma^{-1})$ and $E_{q_i}(\Omega^{-1})$ respectively, where E_{q_i} represents the variational expectation at time i. Following the derivation of the batch updates, our corresponding online learning update for the variational parameters is $q_i(\beta_j) \sim N_{q+1,m}(\hat{\beta}_j^{(i)}, (V_j^{(i)})^{-1} \otimes \Sigma)$, $q_i(\Sigma) \sim IW\left(\nu^{(i)}, S^{(i)}\right)$. $q_i(\tau) \sim IG(a_\tau^{(i)}, b_\tau^{(i)})$, where

$$\hat{\beta}_j^{(i)} = (V_j^{(i)})^{-1}\left((V_j^{(i-1)})\hat{\beta}_j^{(i-1)} + \frac{a_\tau^{(i-1)}}{b_\tau^{(i-1)}}\hat{q}_i(\delta_i = j)E_i y_i^T\right),$$
$$V_j^{(i)} = V_j^{(i-1)} + \frac{a_\tau^{(i-1)}}{b_\tau^{(i-1)}}\hat{q}_i(\delta_i = j)E_i E_i^T,$$

$$\nu^{(i)} = \nu + i.$$
$$S^{(i)} = S^{(i-1)} + \sum_{j=1}^{T}\left\{\frac{a_\tau^{(i-1)}}{b_\tau^{(i-1)}}\hat{q}_i(\delta_i = j)y_i y_i^T + (\hat{\beta}_j^{(i-1)})^T V_j^{(i-1)}\hat{\beta}_j^{(i-1)} - (\hat{\beta}_j^{(i)})^T V_j^{(i)}\hat{\beta}_j^{(i)}\right\},$$
$$a_\tau^{(i)} = a_\tau + \frac{im}{2},$$
$$b_\tau^{(i)} = b_\tau^{(i-1)} + \frac{1}{2}\sum_{j=1}^{T}\hat{q}_i(\delta_i = j)\Big\{(y_i - (\hat{\beta}_j^{(i)})^T E_i)^T \nu^{(i)}(S^{(i)})^{-1}(y_i - (\hat{\beta}_j^{(i)})^T E_i) + m E_i^T (V_j^{(i)})^{-1} E_i)\Big\}.$$

As the second term on the RHS of (8) does not include any terms for $\omega_1, ..., \omega_{N+1}$, there is no online learning required for these parameters. At every step of the online VSUGS algorithm, we continue using the batch update for $q_i(\omega_{1:N+1})$. For the full algorithm, we refer to Algorithm 1.

4.2 Sequential VSUGS type update for the δ_i

Suppose we assimilate observations sequentially and at step $i-1$ we have a variational posterior distribution of the form

$$q_{i-1}(\beta, \Sigma)q_{i-1}(\tau)q_{i-1}(\omega)q_{i-1}(\delta)$$

where $q_{i-1}(\beta,\Sigma) = q_{i-1}(\Sigma)\prod_{j=1}^{T} q(\beta_j|\Sigma)$ with $q_{i-1}(\Sigma)$ being $IW(\nu^{(i-1)}, S^{(i-1)})$, $q_{i-1}(\beta_j|\Sigma)$ being $N_{N+1,m}(\hat{\beta}_j^{(i-1)}, \Omega_j^{(i-1)} \otimes \Sigma)$, $q_{i-1}(\tau) = IG(a_\tau^{(i-1)}, b_\tau^{(i-1)})$ and $q_{i-1}(\omega) = \prod_{j=1}^{N+1} q_{i-1}(\omega_j)$ with $q_{i-1}(\omega_j)$ being $IG(a_j^{(i-1)}, b_j^{(i-1)})$. Also $q_{i-1}(\delta) = \prod_{j=1}^{i-1} q(\delta_j)$. Using the VSUGS approximation, we take

$$\hat{q}_i(\delta_i = j) = r_{ij} \int p(y_i|\beta_{\delta_i}^T E_i, \tau\Sigma) q_{i-1}(\tau) q_{i-1}(\beta_{\delta_i}|\Sigma) q_{i-1}(\Sigma) d\beta_{\delta_i} d\Sigma d\tau \tag{11}$$

where

$$r_{ij} = \begin{cases} \frac{\sum_{k=1}^{i-1} \hat{q}_{i-1}(\delta_k = j) + \alpha/T}{\alpha + i - 1} & j \in \{1, \ldots, \min(i-1, T)\} \\ \frac{\alpha(1 - \min(i-1,T)/T)}{\alpha + i - 1} & j = \min(i-1, T) + 1 \end{cases} . \tag{12}$$

The integral in (12) can be evaluated as (see Appendix B)

$$\begin{aligned} &\int p(y_i|\beta_j^T E_i, \tau\Sigma) q_{i-1}(\tau) q_{i-1}(\beta_j|\Sigma) q_{i-1}(\Sigma) d\beta_j d\Sigma d\tau \\ &= \int (2\pi\tau)^{-\frac{m}{2}} q_{i-1}(\tau) |\Omega_j^{(i-1)}|^{-\frac{m}{2}} |\frac{1}{\tau} E_i E_i^T + (\Omega_j^{(i-1)})^{-1}|^{-\frac{m}{2}} \\ &\quad \frac{|S^{(i-1)}|^{\nu^{(i-1)}/2} \Gamma_m((\nu^{(i-1)}+1)/2)}{2^{\nu^{(i-1)}\frac{m}{2}} \Gamma_m(\nu^{(i-1)}/2)} 2^{(\nu^{(i-1)}+1)\frac{m}{2}} \\ &\quad \left| S^{(i-1)} + \frac{1}{\tau} y_i y_i^T + (\hat{\beta}_j^{(i-1)})^T (\Omega_j^{(i-1)})^{-1} \hat{\beta}_j^{(i-1)} - \bar{\beta}_j^T (\frac{1}{\tau} E_i E_i^T + (\Omega_j^{(i-1)})^{-1}) \bar{\beta}_j \right|^{-\frac{\nu^{(i-1)}+1}{2}} d\tau \end{aligned} \tag{13}$$

where $\bar{\beta}_j = \left(\frac{1}{\tau} E_i E_i^T + (\Omega_j^{(i-1)})^{-1}\right)^{-1} \left(\frac{1}{\tau} E_i y_i^T + (\Omega_j^{(i-1)})^{-1} \hat{\beta}_j^{(i-1)}\right)$. This last integral does not seem to be easily computable analytically. It is an expectation with respect to $q_{i-1}(\tau)$, and if this distribution is concentrated around the mean it is reasonable to make the approximation $\int f(\tau) q_{i-1}(\tau) d\tau = f(E_{q,i-1}(\tau))$ for functions $f(\tau)$ and where we have written $E_{q,i-1}(\tau) = \int \tau q_{i-1}(\tau) d\tau$. Using this approximation here we get that the integral is approximately

$$\begin{aligned} &\left(\frac{2\pi}{\mu_{\tau^{-1}}^{(i-1)}}\right)^{-m/2} |\Omega_j^{(i-1)}|^{-m/2} |\mu_{\tau^{-1}}^{(i-1)} E_i E_i^T + (\Omega_j^{(i-1)})^{-1}|^{-m/2} \\ &\frac{|S^{(i-1)}|^{\nu^{(i-1)}/2} \Gamma_m\left(\frac{\nu^{(i-1)}+1}{2}\right)}{2^{\nu^{(i-1)}m/2} \Gamma_m(\nu^{(i-1)}/2)} 2^{\frac{(\nu^{(i-1)}+1)m}{2}} \\ &\left| S^{(i-1)} + \mu_{\tau^{-1}}^{(i)} y_i y_i^T + (\hat{\beta}_j^{(i-1)})^T (\Omega_j^{(i-1)})^{-1} \hat{\beta}_j^{(i-1)} - \tilde{\beta}_j^T (\mu_{\tau^{-1}}^{(i)} E_i E_i^T + (\Omega_j^{(i-1)})^{-1}) \tilde{\beta}_j \right|^{-\frac{\nu^{(i-1)}+1}{2}} \end{aligned} \tag{14}$$

where $\tilde{\beta}_j = \left(\mu_{\tau^{-1}}^{(i-1)} E_i E_i^T + (\Omega_j^{(i-1)})^{-1}\right)^{-1} \left(\mu_{\tau^{-1}}^{(i-1)} E_i y_i^T + (\Omega_j^{(i-1)})^{-1} \hat{\beta}_j^{(i-1)}\right)$ and $\mu_{\tau^{-1}}^{(i-1)}$ is the expectation of τ^{-1} with respect to $q_{i-1}(\tau)$.

Algorithm 1 : VSUGS algorithm for Matrix DPMM

Initialize $\alpha, T, a_i, b_i, a_\tau, b_\tau, \nu, S, C_1, .., C_T, M_1, .., M_T$

$\hat{q}_1(\delta_1 = 1) \leftarrow 1, \hat{q}_1(\delta_1 = 2), \ldots, \hat{q}_1(\delta_1 = T) \leftarrow 0.$

$a_\tau^{(0)} \leftarrow a_\tau$, $b_\tau^{(0)} \leftarrow b_\tau$, $a_i^{(0)} \leftarrow a_i$, $b_i^{(0)} \leftarrow b_i$, $S^{(0)} \leftarrow S$, $\nu^{(0)} \leftarrow \nu$.

$\mu_{\tau^{-1}}^{(0)} \leftarrow \frac{a_\tau^{(0)}}{b_\tau^{(0)}}$, $\mu_{\Sigma^{-1}}^{(0)} \leftarrow \nu^{(0)}(S^{(0)})^{-1}$, $\mu_{\Omega^{-1}}^{(0)} \leftarrow \text{diag}(a_i^{(0)}/b_i^{(0)})$.

$V_j^{(0)} \leftarrow \mu_{\Omega^{-1}}^{(0)} + C_j$, $\hat{\beta}_j^{(0)} \leftarrow M_j$.

for $i = 1 : n$ **do**

  if $i \geq 2$ **then**

    $T_i \leftarrow \min(T, i-1)$.

    for $j = 1 : T_i$ **do**

      if $j < T_i$ **then**

        $r_{ij} \leftarrow \frac{\sum_{k=1}^{i-1} \hat{q}_{i-1}(\delta_k = j) + \alpha/T}{\alpha + i - 1}$.

      else

        $r_{ij} \leftarrow \frac{\alpha(1 - \min(i-1, T)/T)}{\alpha + i - 1}$.

      end if

      $\Lambda \leftarrow \mu_{\tau^{-1}}^{(i-1)} E_i E_i^T + (\Omega_j^{(i-1)})^{-1}$.

      $\tilde{\beta}_j \leftarrow \Lambda \left(\mu_{\tau^{-1}}^{(i-1)} E_i y_i^T + (\Omega_j^{(i-1)})^{-1} \hat{\beta}_j^{(i-1)} \right)$.

      $\hat{q}_i(\delta_i = j) \leftarrow r_{ij} (\pi(\mu_{\tau^{-1}}^{(i-1)})^{-1})^{-m/2} |\Omega_j^{(i-1)}|^{-m/2} |\Lambda|^{-m/2} \frac{|S^{(i-1)}|^{\nu^{(i-1)}/2}}{\Gamma_m(\nu^{(i-1)}/2)} \Gamma_m \left(\frac{\nu^{(i-1)}+1}{2} \right)$

      $\left| S^{(i-1)} + \mu_{\tau^{-1}}^{(i-1)} y_i y_i^T + (\hat{\beta}_j^{(i-1)})^T (\Omega_j^{(i-1)})^{-1} \hat{\beta}_j^{(i-1)} - \tilde{\beta}_j^T \Lambda \tilde{\beta}_j \right|^{-\frac{\nu^{(i-1)}+1}{2}}$.

    end for

  end if

  for $j = 1 : T$ **do**

    $V_j^{(i)} \leftarrow V_j^{(i-1)} + \mu_{\tau^{-1}}^{(i-1)} \hat{q}_i(\delta_i = j) E_i E_i^T$.

    $\hat{\beta}_j^{(i)} \leftarrow (V_j^{(i)})^{-1} \left(V_j^{(i-1)} \hat{\beta}_j^{(i-1)} + \mu_{\tau^{-1}}^{(i-1)} \hat{q}_i(\delta_i = j) E_i y_i^T \right)$.

  end for

  $\nu^{(i)} \leftarrow \nu^{(i-1)} + 1$.

  $S^{(i)} \leftarrow S^{(i-1)} + \sum_{j=1}^T \left\{ \mu_{\tau^{-1}}^{(i-1)} \hat{q}_i(\delta_i = j) y_i y_i^T + (\hat{\beta}_j^{(i-1)})^T V_j^{(i-1)} \hat{\beta}_j^{(i-1)} - (\hat{\beta}_j^{(i)})^T V_j^{(i)} \hat{\beta}_j^{(i)} \right\}$.

  $a_\tau^{(i)} \leftarrow a_\tau^{(i-1)} + \frac{m}{2}$

  $b_\tau^{(i)} \leftarrow b_\tau^{(i-1)}$

  $\quad + \frac{1}{2} \sum_{j=1}^T \hat{q}_i(\delta_i = j) \left\{ (y_i - (\hat{\beta}_j^{(i)})^T E_i)^T \mu_{\Sigma^{-1}}^{(i)} (y_i - (\hat{\beta}_j^{(i)})^T E_i) + m E_i^T (V_j^{(i)})^{-1} E_i) \right\}$.

  $\mu_{\tau^{-1}}^{(i)} \leftarrow \frac{a_\tau^{(i)}}{b_\tau^{(i)}}$, $\mu_{\Sigma^{-1}}^{(i)} \leftarrow \nu^{(i)}(S^{(i)})^{-1}$.

  for $k = 1 : N + 1$ **do**

    $a_k^{(i)} \leftarrow a_k^{(0)} + \frac{mT}{2}$

    $b_k^{(i)} \leftarrow b_k^{(0)} + \frac{1}{2} \sum_{j=1}^T \left((\hat{\beta}_{j,k}^{(i)} - M_{j,k}^{(i)})^T \mu_{\Sigma^{-1}}^{(i)} (\hat{\beta}_{j,k}^{(i)} - M_{j,k}^{(i)}) + m\omega'_{ij} \right)$,

      where ω_{ij} is the ith diagonal element of $V_j^{(i)}$ and $\hat{\beta}_{j,k}$ is the kth row of $\hat{\beta}_j$.

  end for

end for

5 Posterior predictive inference

Suppose we are given a new input vector x_0 and wish to predict the response vector y_0. Write $E_0 = (1, E_1(x_0), \ldots, E_N(x_0))^T$. The posterior predictive distribution of y_0 can be evaluated as

$$p(y_0|x_0, y_{1:n}) = \sum_{j=1}^{T} p(\delta_0 = j|y_{1:n}) \int p(y_0|y_{1:n}, \tau, \Sigma, \beta_{1:T}, \delta_0 = j) p(\tau, \Sigma, \beta_{1:T}|y_{1:n}) d\beta_{1:T} d\tau d\Sigma.$$

Assuming $n > T$, we replace $p(\delta_0 = j|y_{1:n})$ with $r_{n+1\,j}$ as defined in (12). Also, replacing $p(\tau, \beta_j, \Sigma|y_{1:n})$ with the corresponding variational posterior $q_n(\tau, \beta_j, \Sigma)$, the predictive density becomes

$$p(y_0|E_0, y_{1:n}) = \sum_{j=1}^{T} r_{n+1\,j} \int p(y_0|y_{1:n}, \tau, \Sigma, \beta_j, \delta_0 = j) q_n(\tau) q_n(\beta_j) q_n(\Sigma) d\beta_j d\tau d\Sigma. \quad (15)$$

The integral in (15) evaluates to a multivariate t-distribution. So an approximate posterior predictive density is obtained as a mixture of multivariate t-densities. For more details, including the parameters of the multivariate-t mixture components, see Appendix C.

5.1 Regression-type adjustment for improving predictive inference

One advantage of using the matrix-variate Dirichlet process approach to flexible regression is that avoiding covariate dependence in the mixing weights greatly simplifies computation, something that we have exploited here for implementing an online algorithm. However, this does place a greater burden on the mean functions in the regression mixture components to model the response distribution in a flexible way. Here we consider a method for improving predictive performance of the fitted model, borrowing an idea from the literature on regression adjustment methods for approximate Bayesian computation (Beaumont et al., 2002; Blum, 2010; Blum and François, 2010; Blum and Tran, 2010). The idea below is given in equation (4.1) of Blum and Tran (2010).

Suppose we wish to consider prediction of a new response $y^* = (y_1^*, \ldots, y_m^*)^T$ to be observed with corresponding covariate $x^* = (x_1^*, \ldots, x_p^*)^T$. Write $N_k(x^*)$ for the k nearest neighbours of x^* among the observed covariates $\{x_1, \ldots, x_n\}$. We write the corresponding values of (x, y) as $(x_{i_1}, y_{i_1}), \ldots, (x_{i_k}, y_{i_k})$ so that $i_1, \ldots, i_k$ denote the indices of the covariates in $N_k(x^*)$. In our fitted regression model, write $\hat{F}_j(.|x)$ for the marginal distribution function of the jth component of the response in the fitted model at x, and $\hat{F}_j^{-1}(\cdot|x)$ for

its inverse where it is assumed this exists.

If the fitted model is correct, $\hat{F}_j(y_{i_rj}|x_{ir})$ is uniform on $[0,1]$, and $\hat{F}_j^{-1}(\hat{F}_j(y_{i_rj}|x_{i_r})|x^*)$ has the distribution $\hat{F}_j(\cdot|x^*)$. So if we set

$$\begin{aligned} y_r^a &= (y_{r1}^a, \ldots, y_{rm}^a)^T \\ &= (\hat{F}_1^{-1}(\hat{F}_1(y_{i_r1}|x_{i_r})|x^*), \ldots, \hat{F}_m^{-1}(\hat{F}_m(y_{i_rm}|x_{i_r})|x^*))^T \end{aligned} \tag{16}$$

$r = 1, \ldots, k$, then marginally $y_{1j}^a, \ldots, y_{kj}^a$ is a sample from $\hat{F}_j(\cdot|x^*)$ (if the regression model is correct).

The sample y_r^a, $r = 1, \ldots, k$ can be used to do approximate predictive inference. The advantage of this method compared to using $\hat{F}_j(\cdot|x^*)$ directly is that by using, in effect, quantile residuals locally around x^* to define the particles y_r^a we are able to adjust for any local misfit of the regression model. This can result in improved predictive inference. Note that by transforming the particles y_{i_r} component-wise we are not guaranteed to preserve the correct multivariate dependence structure in the fitted model at x^*, but if the copula of the fitted distribution changes only slowly with x over the neighbourhood used the effects of this approximation are minor.

6 Application to weak informative prior selection

As an application of our proposed methodology, we consider flexible approximation of prior predictive densities as a function of a prior hyperparameter value based on data simulated under a model, when these prior predictive densities are not analytically tractable. Approximating such predictive densities is useful for prior choice. In the application considered here we make use of the way that the MDP mixture model is able to approximate the whole response distribution flexibly. In the next section we will look more closely at the quality of point predictions of the online algorithm compared to those obtained by batch VB and MCMC approaches.

Consider a statistical model $p(y|\xi)$ for data y with parameter ξ. Suppose we have a class of priors $p(\xi|\lambda)$ where λ is a hyperparameter value to be chosen. We also suppose that there is a value λ_0 for λ that has already been chosen tentatively as representing our best current prior knowledge of ξ. For the purpose of sensitivity analysis, we may wish to define a prior that is less informative than $p(\xi|\lambda_0)$, and this might be particularly useful in the case where the information brought by the prior and likelihood seem to be contradictory. Evans and Jang (2011) considered defining the amount of information in a prior $p(\xi|\lambda)$ relative to $p(\xi|\lambda_0)$ through the idea of prior-data conflict. The notion of weakly informative priors formalized in Evans and Jang (2011) was inspired by previous

work of Gelman (2006).

Since the idea of Evans and Jang (2011) is built on the idea of checking for prior-data conflict, this needs to be understood first. Prior-data conflict occurs where the prior puts all its mass out in the tails of the likelihood. A way of testing for prior-data conflict which modifies a suggestion of Box (1980) will be considered here, following Evans and Moshonov (2006). Their idea is that a minimal sufficient statistic value S determines the likelihood, so we can check if the observed likelihood is in conflict with the prior by seeing whether the observed value of the sufficient statistic S_{obs} say lies out in the tails of its prior predictive distribution. A p-value for checking for conflict with the prior $p(\xi|\lambda)$ can be computed as

$$p(S_{obs}, \lambda) = P(p(S|\lambda) \leq p(S_{obs}|\lambda)), \tag{17}$$

where $S \sim p(S|\lambda)$ and $p(S|\lambda) = \int p(S|\xi)p(\xi|\lambda)d\xi$ is the prior predictive distribution of S. If a non-trivial sufficient statistic does not exist it may be reasonable to choose an asymptotically sufficient statistic such as the maximum likelihood estimator or some approximation to it. Note that $p(S_{obs}, \lambda)$ is calculating the probability that a random draw from $p(S|\lambda)$ has lower density than the value of S_{obs} and it is small if S_{obs} lies out in the tails of $p(S|\lambda)$. The above prior-data conflict check can be modified in various ways - for more details see Evans and Moshonov (2006).

To use this notion of prior-data conflict checking to define how informative the prior $p(\xi|\lambda)$ is relative to $p(\xi|\lambda_0)$ Evans and Jang (2011) consider S generated randomly under $p(S|\lambda_0)$ and ask whether for data generated in such a way does doing the analysis under $p(\xi|\lambda)$ rather than $p(\xi|\lambda_0)$ result in a reduction of the frequency of prior-data conflicts. The occurrence of a conflict is defined by choice of a certain cutoff for a conflict p-value such as (17). It is possible to consider various modifications of the basic idea considering uniformity of reduction of levels of conflict over different p-value cutoffs, see Evans and Jang (2011) for more details.

Following the ideas of Evans and Moshonov (2006) and Evans and Jang (2011), Nott et al. (2015) propose modifying a regression adjustment approach used in the approximate Bayesian computation (ABC) literature to approximate prior predictive distributions $p(S|\lambda)$ for many different λ in a computationally thrifty way when S may be expensive to compute. In particular, they consider the method of Blum and François (2010), which modifies a suggestion of Beaumont et al. (2002), to generate approximate samples from prior predictive densities $p(S|\lambda)$ and then use these samples for the required computations. This approach is much more computationally efficient than generating a large number of replications of S at each value of λ independently for every value λ of interest on a

grid, say. The method starts by generating values λ_i, $i = 1, ..., n$ from a pseudo-prior $p(\lambda)$. Then values (ξ_i, S_i), $i = 1, .., n$ are generated for(ξ, S) from $p(\xi|\lambda)p(S|\xi)$ where S is a minimal sufficient statistic or some asymptotically sufficient statistic. Nott et al. (2015) modify the ABC with regression adjustment method in Blum and François (2010) by reversing the usual role of the parameters and the summary statistics where these methods are used in the ABC context. They fit a regression model with

$$S_i = \mu(\lambda_i) + \sigma(\lambda_i)\epsilon_i \tag{18}$$

where the ϵ_i are i.i.d errors with zero mean and variance one and $\mu(\lambda)$ and $\sigma(\lambda)$ are flexible mean and standard deviation functions. Blum and François (2010) parametrize $\mu(\cdot)$ and $\sigma(\cdot)$ using neural networks. After fitting the model to the data to obtain estimates $\hat{\mu}(\lambda)$ and $\hat{\sigma}(\lambda)$, a sample of $p(S|\lambda)$ can be obtained approximately by considering the fitted mean for the regression model plus the empirical residuals. The empirical residual for the ith point is $\hat{\epsilon}_i = \hat{\sigma}(\lambda_i)^{-1}(S_i - \hat{\mu}(\lambda_i))$, and using such empirical residuals together with the fitted model at λ gives

$$S_i^a(\lambda) = \hat{\mu}(\lambda) + \hat{\sigma}(\lambda)\left\{\hat{\sigma}(\lambda_i)^{-1}(S_i - \hat{\mu}(\lambda_i))\right\} \tag{19}$$

as an approximate sample from $p(S|\lambda)$ if the regression model is correct. Based on the approximate sample $S_i^a(\lambda)$, they use a kernel estimate to approximate $p(S|\lambda)$. Let this kernel estimate be $\hat{p}(S|\lambda)$. Next, suppose that S_j^0, $j = 1, \ldots, n$ are draws from $p(S|\lambda_0)$. Then a particle approximation to the distribution of $p(S, \lambda)$ for $S \sim p(S|\lambda_0)$ is given by the values $\hat{P}(S_1^0, \lambda), ..., \hat{P}(S_n^0, \lambda)$ where

$$\hat{P}(S_j^0, \lambda) = \frac{1}{n}\sum_{i=1}^{n} I(\hat{p}(S_i^a(\lambda)|\lambda) \leq \hat{p}(S_j^0|\lambda))$$

The distribution of the p-value can be used to determine whether $p(\xi|\lambda)$ is weakly informative relative to $p(\xi|\lambda_0)$ or not.

We propose using our approach to assess weak informativity of alternative priors compared to a base prior, similar to the above. However, instead of using the ABC with regression adjustment (18), we propose fitting a matrix-variate Dirichlet process mixture model with S as response and λ as predictors. In applying the MDP prior approach we also employ the regression adjustment method of Section 5.1 to obtain approximate samples from $p(S|\lambda)$ at any desired value of λ. Kernel estimates of $p(S|\lambda)$ are then constructed as for the approach of Nott et al. (2015) and the procedure above followed for approximating the distribution of conflict p-values for S generated from $p(S|\lambda_0)$. As

observed in Nott et al. (2015) high accuracy is not needed in the regression calculations; the regression calculations are simply a screening computation, and once a candidate value of λ is chosen for a weakly informative prior then for the single finally chosen value we can generate a large sample from the prior predictive distribution and see whether our approximate calculations were good enough.

6.1 Analysis of a logistic regression example

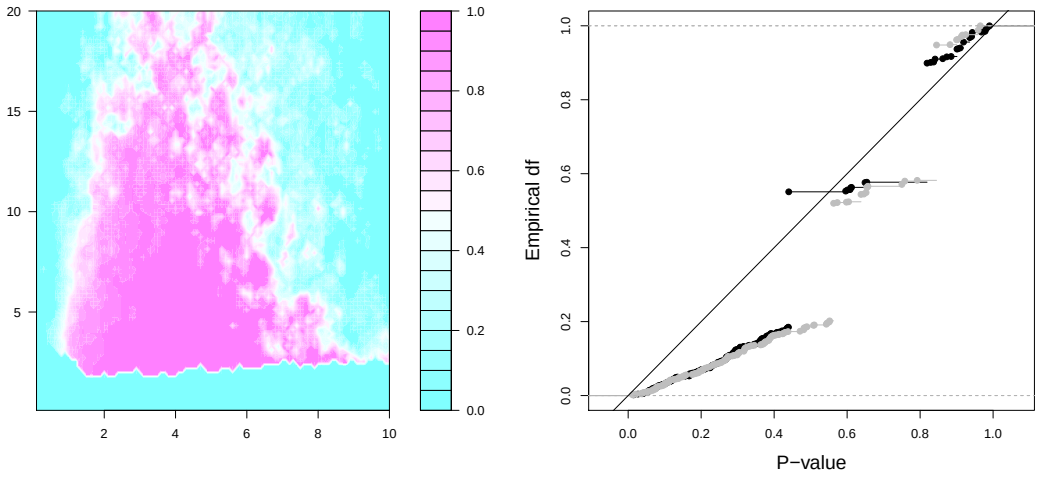

(a) (b)

Figure 1: (a) Estimated degree of weak informativity at $\gamma = 0.05$ for the bioassay example. (b) The plot points in grey and black represent estimated distribution of the conflict p-value using VSUGS-adjusted approach and by direct simulation from the prior predictive distribution. Both estimation uses the alternative prior $\sigma_0 = \sigma_1 = 4$.

We consider a bioassay example from Racine et al. (1986) which is also analysed in Gelman et al. (2008), Evans and Jang (2011) and Nott et al. (2015). In this dataset, four groups of five animals were exposed to different level of doses (x_i) and the number of death (y_i) were recorded. Following Nott et al. (2015) and Evans and Jang (2011), we consider a logistic regression setup. It is assumed that the covariate has been transformed to log scale and centred and scaled as in Gelman et al. (2008). The model is $y_i \sim \text{Bin}(5, p_i)$ where $\text{logit}(p_i) = c_0 + c_1 x_i$. We assume that the priors for c_0 and c_1 are independent and follow Gaussian distributions with zero mean and variances σ_0^2 and σ_1^2 respectively. For our base prior, we consider $\sigma_0 = 10$ and $\sigma_1 = 2.5$.

Evans and Jang (2011) consider the exact sufficient statistics (y_1, y_2, y_3, y_4) for analysis. Nott et al. (2015) consider using the posterior mode $(\hat{c}_0, \hat{c}_1)$ for a prior with $\sigma_0 = \sigma_1 = 10$ as an approximation to the MLE but which unlike the MLE will exist

even in degenerate cases. They consider the MLE for the dimension reduction that it brings and as a generic choice applicable in situations where a non-trivial minimal sufficient statistic doesn't exist. For the statistic S used to define the conflict check in the definition of weak informativity, they use a transformation of $(\hat{c}_0, \hat{c}_1)$ to the fitted probabilities $\hat{p}_2$ and $\hat{p}_3$ at x_2 and x_3 respectively. The reasons for this are discussed further in Nott et al. (2015). That is, our approximate sufficient statistic is $(\hat{p}_2, \hat{p}_3)$ where $\hat{p}_i = 1/(1 + \exp(-\hat{c}_0 + \hat{c}_1 x_i))$ for $i = 2, 3$. Note that, because of the discreteness of the data, strictly the distribution of this statistic is also discrete but continuity may be used as a reasonable approximation when the number of different possible values is large and we do this here. Note also that in the kernel density estimation we ignore any boundary effects due to the bounded support of the statistics.

To use our methodology to investigate weak informativity with respect to the base prior in this example we proceed as follows. First, we generate 400,000 values of (σ_0, σ_1) from a pseudo prior which is uniform distribution on $[0.1, 10] \times [0.1, 20]$. We label these values as $\sigma^{(i)} = (\sigma_0^{(i)}, \sigma_1^{(i)})$ for $i = 1, ..., 400,000$. For each $\sigma^{(i)}$, we generate $c_0^{(i)}$ and $c_1^{(i)}$ from their respectively prior distribution and then compute the probability $(p_1^{(i)}, p_2^{(i)}, p_3^{(i)}, p_4^{(i)})$. We use these probabilities to generate $(y_1^{(i)}, y_2^{(i)}, y_3^{(i)}, y_4^{(i)})$ from their respective binomial distribution.

Let $\hat{p}^{(i)} = (\hat{p}_2^{(i)}, \hat{p}_3^{(i)})^T$ and $K(\sigma^{(i)}, \sigma^{(j)})$ and $K(\sigma^{(i)}, \sigma^{(j)}) = \exp\{-||\sigma^{(i)} - \sigma^{(j)}||/2\kappa^2\}$, where $||.||$ is the Euclidean norm and κ^2 is the mean of the euclidean distance among 5000 random samples drawn from $\sigma^{(1)}, ..., \sigma^{(400,000)}$. Zhang et al. (2010) propose a similar choice of the kernel hyperparameter κ^2 and it is verified to be effective in their experimental analysis. We fit the matrix-variate Dirchlet process mixture model with $\hat{p}^{(i)}$ as our response vector and the basis functions $(1, K(\sigma^{(i)}, \sigma^{(1)}), .., K(\sigma^{(i)}, \sigma^{(N)}))$ as our covariates. We set $\alpha = 100$ and $T = 4$. For our prior, we set $a_\tau = 5$, $b_\tau = 0.5$, $a_i = 5$, $b_i = 0.5$, $S = I_2 + \frac{1}{2}\mathbf{1}_2\mathbf{1}_2^T$, $\nu = 3$ and $M_1, ..., M_T, C_1, .., C_T$ as zero matrices.

We first run Algorithm 1 to fit the model. For each λ_σ on a 100×100 regular grid on $[0.1, 10] \times [0.1, 20]$, we take 1000 nearest neighbours from the set $\{\sigma^{(i)}\}_{i=1,...,400,000}$ using the `knnsearch` function in `matlab`. Then, we use these 1000 nearest neighbour to estimate the corresponding $(\hat{p}_2, \hat{p}_3)^T$ using the regression adjusted approach proposed in Section 5.1. Note that if we were to generate 1000 samples directly for the prior predictive for each of our $10,000$ grid points directly this would increase the number of prior predictive simulations and the computational effort by an order of magnitude. The collection of λ_σ covers the support of the hyperprior and each λ_σ is considered as an alternative prior for comparison with the base prior. We follow the suggestion from Evans and Jang (2011) to measure the degree of weak informativity of an alternative prior. That is, we let p_γ be the $\gamma\%$ quantile of the conflict p-value distribution for the base prior. Under the alternative

prior, let q_γ be the probability of a conflict p-value which is less than or equal to p_γ. The degree of weak informativity ζ_γ is defined as

$$\zeta_\gamma = \begin{cases} 0 & q_\gamma > p_\gamma \\ 1 - q_\gamma/p_\gamma & q_\gamma \leq p_\gamma \end{cases} \tag{20}$$

Figure 1(a) plots the degree of weak informativity ζ_γ for all λ_σ when $\gamma = 0.05$. We observe that the plot is very similar to Figure 2 in Nott et al. (2015). From the plot, it seems that $\sigma_0 = \sigma_1 = 4$ is a suitable choice for a weakly informative prior. This conclusion agrees with the parameter choice of Nott et al. (2015). As mentioned earlier, the approximate regression calculations are simply screening calculations where high accuracy is not needed since the quality of the final answer can be checked. Figure 1(b) shows a comparison of the estimated distribution of the conflict p-values based on regression (grey) compared to one based on direct simulation from the prior predictive at the finally chosen λ (black). In the lower tail, which is what matters for declaring the existence of any conflict and defining weak informativity, the two distributions agree very well, and the approximate regression calculations have successfully allowed us to identify a suitable weakly informative prior. Note that standard procedures such as Gibbs sampler proposed in Zhang et al. (2010) are not suitable for use in this application with the matrix-variate DP prior model as we need to generate a large number of data points from the hyperprior to obtain good estimates of the prior predictive distributions $p(S|\lambda)$ and so a method is needed that is able to handle large datasets.

7 Empirical comparisons of predictive performance

In this section, we focus on the predictive performance of the various approaches when fitting a model. We consider the Gibbs sampler (Zhang et al., 2010) and three different versions of the variational procedure. The first variational approach, which we call VB (Stick breaking), is the method of Blei and Jordan (2006). The second variational method, which we call VB (Pólya urn) is the batch variational method discussed in Section 3. In the third variational approach, we consider the VSUGS approach discussed in Section 4.

In our experiment, the predictors are standardized to have zero mean and unit variance with respect to the training set. We fixed the number of iterations for the Gibbs sampler to be 25000, of which the first 15000 will be discarded as burn in. For the remaining 10000 iterations of the Gibbs sampler, we retain every 10th realization of the parameters. We also fixed the number of iterations for the variational approach using batch updates to 100. For the online variational approach, we first initialize the variational parameters of the assignment variables on a relatively small number of data points using the batch

update and then run Algorithm 1.

We measure the performance of the various approaches by considering their root mean square error (RMSE) and mean absolute percentage error (MAPE). Let $\tilde{y}_1, .., \tilde{y}_m$ be our target response values and $\hat{y}_1, .., \hat{y}_m$ be their respective fitted value. The error indicators are defined as

$$RMSE = \sqrt{\frac{1}{m}\sum_{i=1}^{m}(\tilde{y}_i - \hat{y}_i)^2}$$

and

$$MAPE = \frac{1}{m}\sum_{i=1}^{m}\left|\frac{\tilde{y}_i - \hat{y}_i}{\tilde{y}_i}\right|.$$

Similar performance measures are used in Zhang et al. (2010).

In our analysis, we compute the RMSE and MAPE both in-sample (for the training set) and out-of-sample (for the test set). Although we are mostly interested in out-of-sample predictive performance, looking at in-sample measures of fit can also be useful here where we are comparing several computational approximations for the same posterior; measures of in-sample fit can be revealing about differences in the quality of posterior approximation even if out-of-sample predictive performance is similar for the different methods. For the RMSE and MAPE of in-sample predictions, they are constructed as follows. For the Gibbs sampler, at each retained MCMC realization r, $r = 1, \ldots, 1000$, we have a cluster allocation for each $y_1, ..., y_n$. Suppose let's say that for a particular data point y_i, the allocation at the r realization is the jth component. Then our corresponding regression coefficient estimate would then be $\beta^*_{i,r} = \beta^T_{j,r}$, where $\beta_{1,r}, .., \beta_{T,r}$ are the rth MCMC coefficient matrix realizations. Then, our in-sample fitted value for y_i is estimated as

$$\left\{\sum_{r=1}^{R}\frac{\beta^*_{i,r}}{R}\right\}^T E_i.$$

For all variational procedures, for each y_i, we have the corresponding variational posterior probability of the assignments as well as the posterior mean. We use the weighted sum of these coefficient matrices according to the posterior probability to compute an in-sample fitted value.

Out of sample fitted values are obtained from the posterior predictive distributions of the respective procedures. For the Gibbs sampler, we will use the predictive distribution from equation (7) in Zhang et al. (2010). For the VSUGS procedures, we use the posterior predictive mean to fit each y_i. Details of the posterior predictive mean can be found in Section 5. For prediction accuracy of the test set, we also consider the adjusted VSUGS approach, which is to use the VSUGS with the regression-type adjustment in Section 5.1.

7.1 Energy Data

In this example, we consider the energy efficiency data created by Tsanas and Xifara (2012). This dataset is available at http://archive.ics.uci.edu/ml/datasets.html. Tsanas and Xifara (2012) studied the effect of eight input variables (relative compactness, surface area, wall area, roof area, overall height, orientation, glazing area, glazing area distribution) on two output variables, namely heating load (HL) and cooling load (CL), of residential buildings. The dataset contains 768 instances. We randomly select 100 data points as the test set and use the remaining 668 data points as the training set.

The settings for fitting our model are as follows. We set the number of basis function, N, at 200 and use the same kernel discussed in Section 6.1. Our choice of α is 3, which is set by rounding off the average of the 10000 iterations of α retained from the Gibbs sampler. For comparison of performance, we run both the Gibbs sampler and the variational procedures with a fixed $\alpha = 3$. For the hyperparameters of the priors, we set $a_\tau = 5$, $b_\tau = 0.5$, $a_i = 20$, $b_i = 0.5$, $S = I_2 + \frac{1}{2}\mathbf{1}_2\mathbf{1}_2^T$, $\nu = 3$ and $M_1, ..., M_T, C_1, .., C_T$ as zero matrices. We also set the maximum number of possible components as $T = 10$. For the variational approach with batch update, we initialize each assignment variables randomly to one of the T components and set its variational probability to one. For the matrix VSUGS approach, we initialize the variational parameters for the first 200 assignment variables using the VB(Pólya urn).

Table 1: In-sample accuracy of various approaches for the energy efficiency data.

	RMSE			MAPE			
	y_1	y_2	Mean	y_1	y_2	Mean	Time (mins)
Gibbs Sampler	0.0474	0.0450	0.0462	0.0605	0.0667	0.0636	141
VB (Stick Breaking)	0.1853	0.2150	0.2001	0.2163	0.2469	0.2163	18
VB (Pólya urn)	0.1579	0.1562	0.1570	0.1954	0.2806	0.2380	18
Matrix VSUGS	0.3387	0.3160	0.3273	0.4527	0.5176	0.4851	3

Table 2: Prediction accuracy of various approaches for test set for the energy efficiency data

	RMSE			MAPE		
	y_1	y_2	Mean	y_1	y_2	Mean
Gibbs Sampler	0.5736	0.5881	0.5809	0.7427	0.8055	0.7741
VB (Stick Breaking)	0.3930	0.4911	0.4421	0.5731	0.8347	0.7039
VB (Pólya urn)	0.4005	0.5000	0.4503	0.4784	0.5038	0.4911
Matrix VSUGS	0.4038	0.4882	0.4460	0.5364	0.6140	0.5752
Adjusted VSUGS	0.2558	0.3329	0.2943	0.3399	0.4687	0.4043

Measures of in-sample fit and computation times are presented in Table 2. The two batch VB methods have similar in-sample fits, but the Gibbs sampling and VSUGS

approaches produce quite different results. Table 3 considers out-of-sample predictive accuracy. Table 2 shows that all three variational approaches without the regression-type adjustment perform similarly. The RMSE and MAPE of the Gibbs sampler is higher than the rest. Table 2 also shows that using regression-type adjustment significantly improves prediction accuracy with respect to both RMSE and MAPE.

7.2 Robot Arm Data

In this subsection, we analyse the performance of our proposed algorithm on the robot arm data. This dataset, available from `www.gaussianprocess.org/gpml/data`, relates to an inverse dynamics problem for a seven degrees-of-freedom SARCOS anthropomorphic robot arm. The dataset has 21 covariates and 7 responses and has training and test sets of sizes 44448 and 4449 respectively. The 21 covariates consist of 7 joint positions, 7 joint velocities and 7 joint accelerations and the 7 responses consist of 7 joint torques.

We follow the same procedure used in the energy dataset for our settings. The value of α is set to 12. For the hyperparameters of the prior, we set $a_\tau = 5$, $b_\tau = 0.5$, $a_i = 20$, $b_i = 0.5$, $S = I_7 + \frac{1}{7}\mathbf{1}_7\mathbf{1}_7^T$, $\nu = 8$ and $M_1, ..., M_T, C_1, .., C_T$ as zero matrices. We also set the maximum number of possible components as $T = 10$. For the matrix VSUGS approach, we initialize the variational parameters of the first 500 data points using VB (Pólya urn).

Table 3: In-sample accuracy of various approaches for the robot arm data.

	Method	y_1	y_2	y_3	y_4	y_6	y_6	y_7	Mean
RMSE	Gibbs Sampler	0.1439	0.1338	0.1154	0.0854	0.1615	0.1615	0.1698	0.1304
	VB (Stick Breaking)	0.2058	0.1827	0.1492	0.1512	0.1917	0.2086	0.1521	0.1773
	VB (Pólya urn)	0.1887	0.1976	0.1704	0.1287	0.2042	0.2121	0.1403	0.1774
	Matrix VSUGS	0.5297	0.4665	0.3894	0.4125	0.4531	0.4630	0.4097	0.4463
MAPE	Gibbs Sampler	0.6156	0.4852	0.5554	0.5607	0.5902	0.8916	0.4822	0.5973
	VB (Stick Breaking)	0.8257	0.5441	0.5941	0.7593	0.6417	1.0223	0.5984	0.7122
	VB (Pólya urn)	0.6466	0.5814	0.6625	0.6585	0.6605	0.9861	0.6677	0.6948
	Matrix VSUGS	1.4872	1.0410	1.1798	1.5815	1.1536	1.5569	1.3610	1.3373

The performance of the in-sample fits is presented in Table 3, for fitting to a subsample of size 2000. Corresponding out-of-sample fits for test set subsample of size 500 are shown in Table 4. We also initialize the variational parameters of the first 2000 data points using VB (Pólya urn) and then run Algorithm 1 on the full training dataset of size 44448. As the dataset is large, we do not use the Gibbs sampler or the variational procedures using batch updates for the full dataset. The results of out-of-sample predictive performance are presented in Table 5. It is clear that there is a significant prediction accuracy improvement in terms of RMSE and MAPE when using the adjusted VSUGS.

Table 4: Prediction accuracy of various approaches for the robot arm data.

	Method	y_1	y_2	y_3	y_4	y_6	y_6	y_7	Mean
RMSE	Gibbs Sampler	0.4099	0.3636	0.3404	0.3638	0.3598	0.3867	0.3395	0.3662
	VB (Stick Breaking)	0.5524	0.4909	0.4547	0.4225	0.5227	0.4653	0.4037	0.4732
	VB (Pólya urn)	0.4323	0.4195	0.3983	0.3851	0.4410	0.3916	0.3762	0.4063
	Matrix VSUGS	0.4198	0.3650	0.3375	0.3684	0.3737	0.3867	0.3490	0.3714
	Adjusted VSUGS	0.3513	0.3105	0.2869	0.2910	0.3249	0.3479	0.2741	0.3124
MAPE	Gibbs Sampler	1.0327	0.5987	1.5229	1.4137	1.0385	0.7330	1.2885	1.0897
	VB (Stick Breaking)	1.3711	0.7539	1.9556	1.0920	1.2136	1.1278	1.3825	1.2709
	VB (Pólya urn)	0.7587	0.6397	0.9774	0.9263	1.4217	1.1021	1.1947	1.0029
	Matrix VSUGS	0.9679	0.6600	1.8472	1.4600	1.2228	0.8428	1.4971	1.2140
	Adjusted VSUGS	0.7813	0.6097	1.7723	0.6627	1.1116	0.8206	1.0809	0.9770

Table 5: Prediction accuracy of various approaches using full training set for robot arm data.

Method		y_1	y_2	y_3	y_4	y_6	y_6	y_7	Mean
RMSE	Matrix VSUGS	0.3930	0.3453	0.3065	0.3071	0.3402	0.3330	0.2923	0.3311
	Adjusted VSUGS	0.2505	0.2500	0.2315	0.1626	0.2541	0.2515	0.1785	0.2255
MAPE	Matrix VSUGS	0.8212	0.5288	1.9353	1.0970	0.9967	1.0329	0.7880	1.0285
	Adjusted VSUGS	0.7204	0.4575	1.8183	0.7178	0.8687	0.7910	0.7220	0.8708

7.3 Computation Times

Table 6: Computation time for Robot Arm Data.

Method	Time (minutes)	
Sample size	n = 2000	n = 44448
Gibbs Sampler	2422	-
VB (Stick Breaking)	368	-
VB (Pólya urn)	365	-
Matrix VSUGS	93	584

All the algorithms are run on a `Mac 3.2Ghz i5` Quad core processor with code written in `matlab`. For the energy efficiency dataset, as reflected in Table 1, the computation time for Gibbs sampler, VB (Stick breaking), VB (Pólya urn) and matrix VSUGS require approximately 141, 18, 18 and 3 minutes respectively. For the robot arm dataset, Table 6 shows that the Gibbs sampler, VB (Stick breaking), VB (Pólya urn) and matrix VSUGS require 2422, 368, 365 and 93 minutes respectively. The amount of time required for matrix VSUGS to run the full dataset is 584 minutes, which is significantly shorter than the amount of time it takes for the Gibbs sampler to run on a much smaller dataset. In fact, two third of the computation time is spent on initializing the first 2000 data points.

8 Discussion

In this article, we study variational computational methods for fitting the matrix-variate Dirichlet process mixture model of Zhang et al. (2010), extending the VSUGS approach by Zhang et al. (2014) for Dirichlet process mixtures of normal densities. The method we develop is computationally efficient and especially useful as an alternative to MCMC for analysis of medium to large datasets. In order to increase prediction accuracy, we also propose a regression-type adjustment for improving predictive inference. The adjustment approach is shown to be useful in several real applications.

Appendix A - Variational Batch update

In the derivation of the update for each block γ say we will write simply $E_q(\cdot)$ for the expectation with respect to the current variational posterior distribution q with γ integrated out and will not denote the dependence on the block explicitly in the notation. The meaning will be clear from the context. We consider the mean field update for (β, Σ) first. We have

$$q(\beta,\Sigma) \propto \exp\left(E_q\left(\sum_{i=1}^{n}\sum_{j=1}^{T} I(\delta_i = j)\log p(y_i|\beta_j^T E_i, \tau\Sigma) + \sum_{j=1}^{T}\log p(\beta_j|\Omega,\Sigma) + \log p(\Sigma)\right)\right).$$

Apart from constant terms not depending on β, Σ we have

$$E_q(\log p(\Sigma)) = -\frac{\nu+m+1}{2}\log|\Sigma| - \frac{1}{2}\mathrm{tr}(S\Sigma^{-1})$$

$$E_q(\log p(\beta_j|\Omega,\Sigma)) = -\frac{N+1}{2}\log|\Sigma| - \frac{1}{2}\mathrm{tr}\left((E_q(\Omega^{-1}) + C_j)(\beta_j - M_j)\Sigma^{-1}(\beta_j - M_j)^T\right)$$

$$E_q(\sum_{i=1}^{n}\sum_{j=1}^{T} I(\delta_i = j)\log p(y_i|\beta_j^T E_i, \tau\Sigma)) = \sum_{i=1}^{n}\sum_{j=1}^{T} q_{ij} E_q(\log p(y_i|\beta_j^T E_i, \tau\Sigma)$$

where (again apart from terms not depending on β, Σ)

$$E_q(\log p(y_i|\beta_j^T E_i, \tau\Sigma)) = -\frac{1}{2}\log|\Sigma| - \frac{1}{2}E_q(\tau^{-1})(y_i - \beta_j^T E_i)^T\Sigma^{-1}(y_i - \beta_j^T E_i).$$

This gives (again up to an additive constant)

$$\begin{aligned}\log q(\beta,\Sigma) = &-\frac{(n+N+2)m+\nu+1}{2}\log|\Sigma| - \frac{1}{2}\mathrm{tr}(S\Sigma^{-1}) \\ &- \frac{1}{2}\sum_{j=1}^{T}\mathrm{tr}\left((E_q(\Omega^{-1})+C_j)(\beta_j-M_j)\Sigma^{-1}(\beta_j-M_j)^T\right) \\ &- \frac{1}{2}E_q(\tau^{-1})\sum_{i=1}^{n}\sum_{j=1}^{T}q_{ij}(y_i-\beta_j^TE_i)^T\Sigma^{-1}(y_i-\beta_j^TE_i).\end{aligned}$$

To simplify this, write

$$\hat{\beta}_j = V_j^{-1}\left((E_q(\Omega^{-1})+C_j)M_j + E_q(\tau^{-1})\sum_i q_{ij}E_iy_i^T\right)$$

where

$$V_j = \left(E_q(\Omega^{-1})+C_j+E_q(\tau^{-1})\sum_i q_{ij}E_iE_i^T\right)$$

and observe that

$$\begin{aligned}&\mathrm{tr}\left((E_q(\Omega^{-1})+C_j)(\beta_j-M_j)\Sigma^{-1}(\beta_j-M_j)^T\right) \\ &\quad + E_q(\tau^{-1})\sum_{i=1}^{n}q_{ij}(y_i-\beta_j^TE_i)^T\Sigma^{-1}(y_i-\beta_j^TE_i) \\ &= \mathrm{tr}\left((\beta_j-M_j)^T(E_q(\Omega^{-1})+C_j)(\beta_j-M_j)\Sigma^{-1}\right) \\ &\quad + \mathrm{tr}\left(E_q(\tau^{-1})\sum_{i=1}^{n}q_{ij}(y_i-\beta_j^TE_i)(y_i-\beta_j^TE_i)^T\Sigma^{-1}\right) \\ &= \mathrm{tr}((\beta_j-\hat{\beta}_j)^TV_j(\beta_j-\hat{\beta}_j)\Sigma^{-1}) + \mathrm{tr}\left(M_j^T(E_q(\Omega^{-1})+C_j)M_j\Sigma^{-1}\right) \\ &\quad + \mathrm{tr}\left(E_q(\tau^{-1})\sum_{i=1}^{n}q_{ij}y_iy_i^T\Sigma^{-1}\right) - \mathrm{tr}\left(\hat{\beta}_j^TV_j\hat{\beta}_j\Sigma^{-1}\right).\end{aligned}$$

This means that up to an additive constant

$$\begin{aligned}&\log q(\beta,\Sigma) = \\ &-\frac{(n+N+2)m+\nu+1}{2}\log|\Sigma| - \frac{1}{2}\sum_{j=1}^{T}\mathrm{tr}((\beta_j-\hat{\beta}_j)^TV_j(\beta_j-\hat{\beta}_j)\Sigma^{-1}) \\ &-\frac{1}{2}\mathrm{tr}\left(\left(S+\sum_{j=1}^{T}\left(E_q(\tau^{-1})\sum_{i=1}^{n}q_{ij}y_iy_i^T + M_j^T(E_q(\Omega^{-1})+C_j)M_j - \hat{\beta}_j^TV_j\hat{\beta}_j\right)\right)\Sigma^{-1}\right).\end{aligned}$$

Then we recognize the form of $q(\beta, \Sigma)$ as being $q(\beta, \Sigma) = q(\Sigma)q(\beta|\Sigma)$ where $q(\Sigma)$ is inverse Wishart,

$$IW\left(\nu + n, S + \sum_{j=1}^{T}\left(E_q(\tau^{-1})\sum_{i=1}^{n} q_{ij} y_i y_i^T + M_j^T(E_q(\Omega^{-1}) + C_j)M_j - \hat{\beta}_j^T V_j \hat{\beta}_j\right)\right)$$

and $q(\beta|\Sigma) = \prod_{j=1}^{T} q(\beta_j|\Sigma)$ with $q(\beta_j|\Sigma) = N_{N+1,m}(\hat{\beta}_j, V_j^{-1} \otimes \Sigma)$.

Next, let's consider the mean field update for $q(\tau)$. We have

$$q(\tau) \propto \exp(E_q(\sum_{i=1}^{n}\sum_{j=1}^{T} I(\delta_i = j)\log p(y_i|\beta_j^T E_i, \tau\Sigma) + \log p(\tau)).$$

Apart from additive constants

$$E_q(\log p(\tau)) = \log p(\tau) = -(a_\tau + 1)\log\tau - b_\tau/\tau$$

and

$$E_q\left(\sum_{i=1}^{n}\sum_{j=1}^{T} I(\delta_i = j)\log p(y_i|\beta_j^T E_i, \tau\Sigma)\right) = \sum_{i=1}^{n}\sum_{j=1}^{T} q_{ij} E_q(\log p(y_i|\beta_j^T E_i, \tau\Sigma)$$

where

$$E_q(\log p(y_i|\beta_j^T E_i, \tau\Sigma)) = \frac{-nm}{2}\log\tau - \frac{1}{2\tau}E_q((y_i - \beta_j^T E_i)\Sigma^{-1}(y_i - \beta_j^T E_i)).$$

Next,

$$E_q((y_i - \beta_j^T E_i)^T\Sigma^{-1}(y_i - \beta_j^T E_i)) = E_q(E_q((y_i - \beta_j^T E_i)^T\Sigma^{-1}(y_i - \beta_j^T E_i)|\Sigma)).$$

To evaluate the inner conditional expectation, we use the following Lemma (Guptar and Nagar, p. 60).

Lemma 1. *Suppose that $X \sim N_{p,q}(M, \Delta \otimes \psi)$. Let A be a $p \times p$ matrix. Then*

$$E(X^T AX) = tr(\Delta A^T)\psi + M^T AM$$

Using Lemma 1 and some simple algebra we obtain

$$E_q((y_i - \beta_j^T E_i)^T\Sigma^{-1}(y_i - \beta_j^T E_i)|\Sigma) = (y_i - \hat{\beta}_j^T E_i)^T\Sigma^{-1}(y_i - \hat{\beta}_j^T E_i) + mE_i^T V_j^{-1} E_i.$$

Hence

$$E_q((y_i - \beta_j^T E_i)^T \Sigma^{-1}(y_i - \beta_j^T E_i)) = (y_i - \hat{\beta}_j^T E_i)^T E_q(\Sigma^{-1})(y_i - \hat{\beta}_j^T E_i) + m E_i^T V_j^{-1} E_i.$$

So apart from additive constants

$$\log q(\tau) = -(a_\tau + \frac{nm}{2} + 1)\log\tau \\ - \frac{1}{\tau}\left\{ b_\tau + \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{T} q_{ij}((y_i - \hat{\beta}_j^T E_i)^T E_q(\Sigma^{-1})(y_i - \hat{\beta}_j^T E_i) + m E_i^T V_j^{-1} E_i)) \right\}.$$

Hence we recognize that $q(\tau)$ is inverse gamma,

$$IG(a_\tau + \frac{nm}{2}, b_\tau + \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{T} q_{ij}((y_i - \hat{\beta}_j^T E_i)^T E_q(\Sigma^{-1})(y_i - \hat{\beta}_j^T E_i) + m E_i^T V_j^{-1} E_i)).$$

Next we consider the variational update for ω. We have that apart from additive constants

$$\log q(\omega) = -\sum_{i=1}^{N+1}(a_i + 1)\log\omega_i - \sum_{i=1}^{N+1}\frac{b_i}{\omega_i} - \frac{mT}{2}\sum_{i=1}^{N+1}\log\omega_i \\ - \frac{1}{2}\sum_{i=1}^{N+1}\frac{1}{\omega_i}\sum_{j=1}^{T} E_q((\beta_{j,i} - M_{j,i})^T \Sigma^{-1}(\beta_{j,i} - M_{j,i}))$$

where $\beta_{j,i}$ and $M_{j,i}$ denote the ith rows of β_j and M_j respectively. Writing

$$E_q((\beta_{j,i} - M_{j,i})^T \Sigma^{-1}(\beta_{j,i})) = E_q(E_q((\beta_{j,i} - M_{j,i})^T \Sigma^{-1}(\beta_{j,i})|\Sigma))$$

and noting that $\beta_{j,i}|\Sigma \sim N(\hat{\beta}_{j,i}, \omega'_{ij}\Sigma)$ where $\hat{\beta}_{j,i}$ is the ith row of $\hat{\beta}_j$ and ω'_{ij} is the ith diagonal element of V_j^{-1} we have

$$\begin{aligned} E_q((\beta_{j,i} - M_{j,i})^T \Sigma^{-1}(\beta_{j,i})) &= (\hat{\beta}_{j,i} - M_{j,i})^T \Sigma^{-1}(\hat{\beta}_{j,i} - M_{j,i}) + \mathrm{tr}(\Sigma^{-1}\omega'_{ij}\Sigma) \\ &= (\hat{\beta}_{j,i} - M_{j,i})^T E_q(\Sigma^{-1})(\hat{\beta}_{j,i} - M_{j,i}) + \omega'_{ij} m. \end{aligned}$$

So apart from additive constants

$$\log q(\omega) = -\sum_{i=1}^{N+1}(a_i + mT/2 + 1)\log\omega_i \\ - \sum_{i=1}^{N+1}\frac{1}{\omega_i}(b_i + \frac{1}{2}\sum_{j=1}^{T}\left((\hat{\beta}_{j,i} - M_{j,i})^T E_q(\Sigma^{-1})(\hat{\beta}_{j,i} - M_{j,i}) + \omega'_{ij} m\right)$$

and so we recognize that $q(\omega) = \prod_{i=1}^{N+1} q(\omega_i)$ where $q(\omega_i)$ is inverse gamma,

$$IG(a_i + mT/2, b_i + \frac{1}{2}\sum_{j=1}^{T}\Big((\hat{\beta}_{j,i} - M_{j,i})E_q(\Sigma^{-1})(\hat{\beta}_{j,i} - M_{j,i})^T + m\omega'_{ij}\Big)). \tag{21}$$

Appendix B - Sequential update for δ_i

To help evaluate the integral we use the following lemma.

Lemma 2. *(a) Let Z, V, W and C be matrices with Z $s \times t$, $V > 0$ $s \times s$, $W > 0$ $t \times t$ and C $s \times t$. Then*

$$\int \exp(tr(-\frac{1}{2}(V^{-1}ZW^{-1}Z^T - 2V^{-1}CW^{-1}Z^T))dZ$$
$$= (2\pi)^{st/2}|V|^{t/2}|W|^{s/2}\exp(tr(\frac{1}{2}V^{-1}CW^{-1}C^T)).$$

(b) Suppose ψ, A are matrices with $\psi > 0$ $m \times m$ and $A > 0$ $m \times m$. Let $a > 0$ be a constant. Then

$$\int |\psi|^{-(a+m+1)/2}\exp(tr(-\frac{1}{2}A\psi^{-1}))d\psi = |A|^{-a/2}2^{am/2}\Gamma_m(a/2)$$

where $\Gamma_m(x)$ is the multivariate gamma function, $\Gamma_m(x) = \pi^{m(m-1)/4}\prod_{i=1}^{m}\Gamma(x + (1-i)/2)$.

Proof. Part a) of the lemma follows easily from the fact that the integrand is an unnormalized matrix normal distribution. Using the fact that the integral of the corresponding normalized density is one and rearranging gives the result. Part b) follows in a similar way by noting that the integrand is an unnormalized inverse Wishart density. □

Using Lemma 2 to help evaluate the integral we get

$$\int p(y_i|\beta_j^T E_i, \tau\Sigma)q_{i-1}(\tau)q_{i-1}(\beta_j|\Sigma)q_{i-1}(\Sigma)d\beta_j d\Sigma d\tau$$
$$= \int\int (2\pi\tau)^{-m/2}|\Sigma|^{-1/2}q_{i-1}(\tau)q_{i-1}(\Sigma)\exp\left\{-\frac{1}{2\tau}y_i^T\Sigma^{-1}y_i\right\}(2\pi)^{-(N+1)m/2}|$$
$$\Omega_j^{(i-1)}|^{-m/2}|\Sigma|^{-(N+1)/2}\exp(\mathrm{tr}(-\frac{1}{2}(\Omega_j^{(i-1)})^{-1}\hat{\beta}_j^{(i-1)}\Sigma^{-1}(\hat{\beta}_j^{(i-1)})^T))$$
$$\int \exp\left\{\mathrm{tr}\left[-\frac{1}{2}\left((\frac{1}{\tau}E_iE_i^T + (\Omega_j^{(i-1)})^{-1})\beta_j\Sigma^{-1}\beta_j^T\right.\right.\right.$$
$$\left.\left.\left. -2\left(\frac{1}{\tau}E_iE_i^T + (\Omega_j^{(i-1)})^{-1}\right)\bar{\beta}_j\Sigma^{-1}\beta_j^T\right)\right]\right\}d\beta_j d\Sigma d\tau$$

where $\bar{\beta}_j = \left(\frac{1}{\tau}E_iE_i^T + (\Omega_j^{(i-1)})^{-1}\right)^{-1}\left(\frac{1}{\tau}E_iy_i^T + (\Omega_j^{(i-1)})^{-1}\hat{\beta}_j^{(i-1)}\right)$. Next, use part (a) of Lemma 2 and the explicit form for $q_{i-1}(\Sigma)$ to get

$$\int (2\pi\tau)^{-m/2} q_{i-1}(\tau)|\Omega_j^{(i-1)}|^{-m/2}|\frac{1}{\tau}E_iE_i^T + (\Omega_j^{(i-1)})^{-1}|^{-m/2}\frac{|S^{(i-1)}|^{\nu^{(i-1)}/2}}{2^{\nu^{(i-1)}m/2}\Gamma_m(\nu^{(i-1)}/2)}$$
$$\int |\Sigma|^{-(\nu^{(i-1)}+m+2)/2}\exp(\mathrm{tr}(-\frac{1}{2}(S^{(i-1)} + \frac{1}{\tau}y_iy_i^T + (\hat{\beta}_j^{(i-1)})^T(\Omega_j^{(i-1)})^{-1}\hat{\beta}_j^{(i-1)}$$
$$- \bar{\beta}_j^T(\frac{1}{\tau}E_iE_i^T + (\Omega_j^{(i-1)})^{-1})\bar{\beta}_j)\Sigma^{-1}))d\Sigma d\tau.$$

Using part (b) of Lemma 2, this is equal to

$$\int (2\pi\tau)^{-\frac{m}{2}} q_{i-1}(\tau)|\Omega_j^{(i-1)}|^{-\frac{m}{2}}|\frac{1}{\tau}E_iE_i^T + (\Omega_j^{(i-1)})^{-1}|^{-\frac{m}{2}}$$
$$\frac{|S^{(i-1)}|^{\nu^{(i-1)}/2}\Gamma_m((\nu^{(i-1)}+1)/2)}{2^{\nu^{(i-1)}\frac{m}{2}}\Gamma_m(\nu^{(i-1)}/2)}2^{(\nu^{(i-1)}+1)\frac{m}{2}}$$
$$\left|S^{(i-1)} + \frac{1}{\tau}y_iy_i^T + (\hat{\beta}_j^{(i-1)})^T(\Omega_j^{(i-1)})^{-1}\hat{\beta}_j^{(i-1)} - \bar{\beta}_j^T(\frac{1}{\tau}E_iE_i^T + (\Omega_j^{(i-1)})^{-1})\bar{\beta}_j\right|^{-\frac{\nu^{(i-1)}+1}{2}} d\tau.$$

Appendix C - Posterior predictive distribution

We now evaluate the integral in (15). Following the same steps to get (14), we have

$$\int p(y_0|y_{1:n},\tau,\Sigma,\beta_j,\delta_0=j)q_n(\tau)q_n(\beta_j)q_n(\Sigma)d\beta_j d\tau d\Sigma$$
$$= \left(\frac{2\pi}{\mu_{\tau^{-1}}^{(n)}}\right)^{-m/2}|\Omega_j^{(n)}|^{-m/2}|\mu_{\tau^{-1}}^{(n)}E_0E_0^T + (\Omega_j^{(n)})^{-1}|^{-m/2}\frac{2^{\frac{(\nu^{(n)}+1)m}{2}}\Gamma_m\left(\frac{\nu^{(n)}+1}{2}\right)}{2^{\nu^{(n)}m/2}\Gamma_m(\nu^{(n)}/2)}|S^{(n)}|^{\nu^{(n)}/2}$$
$$|S^{(n)} + \mu_{\tau^{-1}}^{(n)}y_0y_0^T + (\hat{\beta}_j^{(n)})^T(\Omega_j^{(n)})^{-1}\hat{\beta}_j^{(n)} - \tilde{\beta}_j^T(\mu_{\tau^{-1}}^{(n)}E_0E_0^T + (\Omega_j^{(n)})^{-1})\tilde{\beta}_j|^{-\frac{\nu^{(n)}+1}{2}}.$$

where $\tilde{\beta}_j = (\mu_{\tau^{-1}}^{(n)}E_0E_0^T + (\Omega_j^{(n)})^{-1})^{-1}(\mu_{\tau^{-1}}^{(n)}E_0y_0^T + (\Omega_j^{(n)})^{-1}\hat{\beta}_j^{(n)})$. We now show that after simplification of the term inside the determinant we obtain the multivariate t-distribution. We require here the matrix determinant lemma, which states the following.

Lemma 3. *For any invertible matrix A and vector u and v we have*

$$|A + uv^T| = (1 + v^TA^{-1}u)|A|.$$

Letting $\Lambda = (\mu_{\tau^{-1}}^{(n)} E_0 E_0^T + (\Omega_j^{(n)})^{-1})$, it is clear that

$$\begin{aligned}
&\tilde{\beta}_j^T \left(\mu_{\tau^{-1}}^{(n)} E_0 E_0^T + (\Omega_j^{(n)})^{-1}\right) \tilde{\beta}_j \\
&= \left(\mu_{\tau^{-1}}^{(n)} E_0 y_0^T + (\Omega_j^{(n)})^{-1} \hat{\beta}_j^{(n)}\right)^T \Lambda^{-1} \Lambda \Lambda^{-1} \left(\mu_{\tau^{-1}}^{(n)} E_0 y_0^T + (\Omega_j^{(n)})^{-1} \hat{\beta}_j^{(n)}\right) \\
&= (\mu_{\tau^{-1}}^{(n)})^2 y_0 E_0^T \Lambda^{-1} E_0 y_0^T + \mu_{\tau^{-1}}^{(n)} y_0 E_0^T \Lambda^{-1} (\Omega_j^{(n)})^{-1} \hat{\beta}_j^{(n)} \\
&\quad + \mu_{\tau^{-1}}^{(n)} (\hat{\beta}_j^{(n)})^T (\Omega_j^{(n)})^T \Lambda^{-1} E_0 y_0^T + (\hat{\beta}_j^{(n)})^T (\Omega_j^{(n)})^T \Lambda^{-1} \Omega_j^{(n)} \hat{\beta}_j^{(n)}.
\end{aligned}$$

This gives us

$$\begin{aligned}
&\left| S^{(n)} + \mu_{\tau^{-1}}^{(n)} y_0 y_0^T + (\hat{\beta}_j^{(n)})^T (\Omega_j^{(n)})^{-1} \hat{\beta}_j^{(n)} - \tilde{\beta}_j^T (\mu_{\tau^{-1}}^{(n)} E_0 E_0^T + (\Omega_j^{(n)})^{-1}) \tilde{\beta}_j \right|^{-\frac{\nu^{(n)}+1}{2}} \\
&= \Bigg| S^{(n)} + \mu_{\tau^{-1}}^{(n)} y_0 y_0^T + (\hat{\beta}_j^{(n)})^T (\Omega_j^{(n)})^{-1} \hat{\beta}_j^{(n)} - (\mu_{\tau^{-1}}^{(n)})^2 y_0 E_0^T \Lambda^{-1} E_0 y_0^T \\
&\quad - \mu_{\tau^{-1}}^{(n)} y_0 E_0^T \Lambda^{-1} \Omega_j^{(n)} \hat{\beta}_j^{(n)} - \mu_{\tau^{-1}}^{(n)} (\hat{\beta}_j^{(n)})^T (\Omega_j^{(n)})^T \Lambda^{-1} E_0 y_0^T \\
&\quad - (\hat{\beta}_j^{(n)})^T (\Omega_j^{(n)})^T \Lambda^{-1} \Omega_j^{(n)} \hat{\beta}_j^{(n)} \Bigg|^{-\frac{\nu^{(n)}+1}{2}} \\
&= \Bigg| S^{(n)} + \mu_{\tau^{-1}}^{(n)} y_0 (I - \mu_{\tau^{-1}}^{(n)} E_0^T \Lambda^{-1} E_0) y_0^T - \mu_{\tau^{-1}}^{(n)} y_0 E_0^T \Lambda^{-1} \Omega_j^{(n)} \hat{\beta}_j^{(n)} \\
&\quad - \mu_{\tau^{-1}}^{(n)} (\hat{\beta}_j^{(n)})^T (\Omega_j^{(n)})^T \Lambda^{-1} E_0 y_0^T - (\hat{\beta}_j^{(n)})^T (\Omega_j^{(n)})^T \Lambda^{-1} \Omega_j^{(n)} \hat{\beta}_j^{(n)} \\
&\quad + (\hat{\beta}_j^{(n)})^T (\Omega_j^{(n)})^{-1} \hat{\beta}_j^{(n)} \Bigg|^{-\frac{\nu^{(n)}+1}{2}}.
\end{aligned}$$

Let $S_* = S^{(n)} - (\hat{\beta}_j^{(n)})^T (\Omega_j^{(n)})^T \Lambda^{-1} \Omega_j^{(n)} \hat{\beta}_j^{(n)} + (\hat{\beta}_j^{(n)})^T (\Omega_j^{(n)})^{-1} \hat{\beta}_j^{(n)}$ (where we suppress dependence on j in the notation). Using Lemma 3, we have

$$\left| 1 + \mu_{\tau^{-1}}^{(n)} y_0^T S_*^{-1} y_0 (1 - \mu_{\tau^{-1}}^{(n)} E_0^T \Lambda^{-1} E_0) - 2\mu_{\tau^{-1}}^{(n)} E_0^T \Lambda^{-1} \Omega_j^{(n)} \hat{\beta}_j^{(n)} S_*^{-1} y_0 \right|^{-\frac{\nu^{(n)}+1}{2}} |S_*|^{-\frac{\nu^{(n)}+1}{2}}.$$

As the term inside the absolute value is a quadratic function in y_0, the predictive distribution is a multivariate t-distribution. Letting $\mathcal{A} = \mu_{\tau^{-1}}^{(n)} (1 - \mu_{\tau^{-1}}^{(n)} E_0^T \Lambda^{-1} E_0) S_*^{-1}$, $\mathcal{B} = -2\mu_{\tau^{-1}}^{(n)} S_*^{-1} (\hat{\beta}_j^{(n)})^T \Omega_j^{(n)} \Lambda^{-1} E_0$ (again suppressing dependence on j in the notation),

we have

$$\begin{aligned}
&\left|1+\mu_{\tau^{-1}}^{(n)} y_0^T S_*^{-1} y_0(1-\mu_{\tau^{-1}}^{(n)} E_0^T \Lambda^{-1} E_0) - 2\mu_{\tau^{-1}}^{(n)} y_0^T S_*^{-1} (\hat{\beta}_j^{(n)})^T \Omega_j^{(n)} \Lambda^{-1} E_0\right| \\
&= \left|1+(y_0+\frac{1}{2}\mathcal{A}^{-1}\mathcal{B})^T \mathcal{A}(y_0+\frac{1}{2}\mathcal{A}^{-1}\mathcal{B}) - \frac{1}{4}\mathcal{B}^T\mathcal{A}^{-1}\mathcal{B}\right| \\
&= \left|1+(y_0+\frac{1}{2}\mathcal{A}^{-1}\mathcal{B})^T \mathcal{A}(1-\frac{1}{4}\mathcal{B}^T\mathcal{A}^{-1}\mathcal{B})^{-1}(y_0+\frac{1}{2}\mathcal{A}^{-1}\mathcal{B})\right| \left|1-\frac{1}{4}\mathcal{B}^T\mathcal{A}^{-1}\mathcal{B}\right|.
\end{aligned}$$

Thus, the jth component in the expression for the posterior predictive density (15) is a multivariate t-density with location $-\frac{1}{2}\mathcal{A}^{-1}\mathcal{B}$ and variance $\frac{1}{\nu^{(n)}-m-1}\left\{\mathcal{A}(1-\frac{1}{4}\mathcal{B}^T\mathcal{A}^{-1}\mathcal{B})^{-1}\right\}^{-1}$. Hence (15) is approximately a mixture of multivariate t-densities.

Appendix D - Variational Lower bound

Next we compute the variational lower bound on $\log p(y_{1:n})$, which is defined as

$$\mathcal{L}(q) = E_q\left\{\log\left[p\left(y_{1:n}, \beta_{1:T}, \delta_{1:n}, \tau, \omega_{1:(N+1)}, \Sigma\right)\right]\right\} - E_q\left\{\log\left[q\left(\beta_{1:T}, \delta_{1:n}, \tau, \omega_{1:(N+1)}, \Sigma\right)\right]\right\}.$$

Similar to the approach used to calculate the lower bound for the normal mixture model in Zhang et al. (2014), we approximate $\mathcal{L}(q)$ recursively. Let $\theta_t = (\beta_t, \tau, \omega_{1:N+1}, \Sigma)$ and using r_{ij} as an approximation to $p(\delta_i|y_{i:i-1})$, we have

$$\begin{aligned}
\mathcal{L}(q) &= E_{q_i}\left\{\log(q_{i-1}(\beta_{1:T}, \tau, \omega_{1:(N+1)}, \Sigma)p(\delta_i|y_{i:i-1})p(y_i|\delta_i, \theta))\right\} - E_{q_i}\left\{\log(q_i(\theta)q_i(\delta_i))\right\} \\
&= E_{q_i}\Bigg\{\log(q_{i-1}(\tau)) + \log(q_{i-1}(\Sigma)) + \sum_{j=1}^{N+1}\log(q_{i-1}(\omega_j) + \sum_{t=1}^{T}\log(q_{i-1}(\beta_t)) + \log(r_{ij}) + \\
&\quad \log(y_i|\theta_{\delta_i})\Bigg\} - E_{q_i}\Bigg\{q_i(\tau) + \log(q_i(\Sigma)) + \sum_{j=1}^{N+1}\log(q_i(\omega_j) \\
&\quad + \sum_{t=1}^{T}\log(q_i(\beta_t)) + \log(q_i(\delta_i))\Bigg\}.
\end{aligned}$$

Suppressing the expectations with respect to q_i, evaluating the terms involving τ and ω_j gives us

$$\begin{aligned}
E_{q_i}\left\{\log(q_{i-1}(\tau))\right\} - E_{q_i}\left\{\log(q_i(\tau))\right\} &= (a_\tau^{(i)} - a_\tau^{(i-1)})\psi(a_\tau^{(i)}) - \log(\Gamma(a_\tau^{(i)})) + \log(\Gamma(a_\tau^{(i-1)})) \\
&\quad + a_\tau^{(i-1)}(\log(b_\tau^{(i-1)}) - \log(b_\tau^{(i)})) + a_\tau^{(i)}\frac{b_\tau^{(i)} - b_\tau^{(i-1)}}{b_\tau^{(i)}},
\end{aligned}$$

$$E_{q_i}\left\{\log(q_{i-1}(\omega_j))\right\} - E_{q_i}\left\{\log(q_i(\omega_j))\right\} = (a_j^{(i)} - a_j^{(i-1)})\psi(a_j^{(i)}) - \log(\Gamma(a_j^{(i)})) + \log(\Gamma(a_j^{(i-1)})) + a_j^{(i-1)}(\log(b_j^{(i-1)}) - \log(b_j^{(i)})) + a_j^{(i)}\frac{b_j^{(i)} - b_j^{(i-1)}}{b_j^{(i)}}.$$

Since $q_i(\Sigma)$ follows an inverse Wishart distribution with degrees of freedom $v^{(i)}$ and scale matrix $S^{(i)}$, we obtain

$$\begin{aligned} E_{q_i}\left\{\log(q_{i-1}(\Sigma))\right\} &= \frac{\nu^{(i-1)}}{2}\log|S^{(i-1)}| - \frac{\nu^{(i-1)}m}{2}\log(2) - \log\Gamma_m(\frac{\nu^{(i-1)}}{2}) \\ &\quad - \frac{1}{2}tr\left\{S^{(i-1)}E_{q_i}(\Sigma^{-1})\right\} - \frac{\nu^{(i-1)}+m+1}{2}E_{q_i}\left(\log|\Sigma|\right) \\ &= \frac{\nu^{(i-1)}}{2}\log|S^{(i-1)}| - \frac{\nu^{(i-1)}m}{2}\log(2) - \log\Gamma_m(\frac{\nu^{(i-1)}}{2}) \\ &\quad - \frac{1}{2}tr\left\{S^{(i-1)}(S^{(i)})^{-1}\nu^{(i)}\right\} \\ &\quad - \frac{\nu^{(i-1)}+m+1}{2}\left\{-\psi_m\left(\frac{\nu^{(i)}}{2}\right) - m\log(2) + \log|S^{(i)}|\right\} \end{aligned}$$

$$\begin{aligned} E_{q_i}\left\{\log(q_i(\Sigma))\right\} &= \frac{\nu^{(i)}}{2}\log|S^{(i)}| - \frac{\nu^{(i)}m}{2}\log(2) - \log\Gamma_m(\frac{\nu^{(i)}}{2}) - \frac{1}{2}tr\left\{S^{(i)}E_{q_i}(\Sigma^{-1})\right\} \\ &\quad - \frac{\nu^{(i)}+m+1}{2}E_{q_i}\left(\log|\Sigma|\right) \\ &= \frac{\nu^{(i)}}{2}\log|S^{(i)}| - \frac{\nu^{(i)}m}{2}\log(2) - \log\Gamma_m(\frac{\nu^{(i)}}{2}) - \frac{1}{2}tr\left\{S^{(i)}(S^{(i)})^{-1}\nu^{(i)}\right\} \\ &\quad - \frac{\nu^{(i)}+m+1}{2}\left\{-\psi_m\left(\frac{\nu^{(i)}}{2}\right) - m\log(2) + \log|S^{(i)}|\right\}. \end{aligned}$$

Next, observe that if Z is matrix-variate normal distribution $N_{s,t}(C, V \otimes W)$, then $E(Z^TBZ) = W\mathrm{tr}(VB^T) + C^TBC$. Moreover, for A a $r \times s$ matrix, we have $AX \sim N_{r,t}(AC, AVA^T \otimes W)$ and $E(AX\Lambda A^TX^T) = AVA^T\mathrm{tr}(\Lambda W) + (AC)^T\Lambda AC$. Therefore, we

have

$$
\begin{aligned}
E_{q_i}\left\{\log(q_{i-1}(\beta_j))\right\} &= -\frac{(N+1)m}{2}\log(2\pi) - \frac{N+1}{2}E_{q_i}(\log|\Sigma|) - \frac{m}{2}\log|(V_j^{(i-1)})^{-1}| \\
&\quad - E_{q_i}\left\{\frac{1}{2}\mathrm{tr}\left[\Sigma^{-1}(\beta_j - \hat{\beta}_j^{(i-1)})^T V_j^{(i-1)}(\beta_j - \hat{\beta}_j^{(i-1)})\right]\right\} \\
&= -\frac{(N+1)m}{2}\log(2\pi) - \frac{N+1}{2}E_{q_i}(\log|\Sigma|) - \frac{m}{2}\log|(V_j^{(i-1)})^{-1}| \\
&\quad - \frac{1}{2}\mathrm{tr}\Bigg[\nu^{(i)}(S^{(i)})^{-1}(\hat{\beta}_j^{(i)} - \hat{\beta}_j^{(i-1)})^T V_j^{(i-1)}(\hat{\beta}_j^{(i)} - \hat{\beta}_j^{(i-1)}) \\
&\quad + E_{q_i}\left(\Sigma^{-1}\Sigma\mathrm{tr}((V_j^{(i)})^{-1}V_j^{(i-1)})\right)\Bigg] \\
&= -\frac{(N+1)m}{2}\log(2\pi) - \frac{N+1}{2}E_{q_i}(\log|\Sigma|) - \frac{m}{2}\log|(V_j^{(i-1)})^{-1}| \\
&\quad - \frac{1}{2}\mathrm{tr}\Bigg[\nu^{(i)}(S^{(i)})^{-1}(\hat{\beta}_j^{(i)} - \hat{\beta}_j^{(i-1)})^T V_j^{(i-1)}(\hat{\beta}_j^{(i)} - \hat{\beta}_j^{(i-1)}) \\
&\quad + I_m\mathrm{tr}((V_j^{(i)})^{-1}V_j^{(i-1)})\Bigg]
\end{aligned}
$$

$$
\begin{aligned}
E_{q_i}\left\{\log(q_i(\beta_j))\right\} &= -\frac{(N+1)m}{2}\log(2\pi) - \frac{N+1}{2}E_{q_i}(\log|\Sigma|) - \frac{m}{2}\log|(V_j^{(i)})^{-1}| \\
&\quad - E_{q_i}\left\{\frac{1}{2}\mathrm{tr}\left[(V_j^{(i)})(\beta_j - \hat{\beta}_j^{(i)})^T\Sigma^{-1}(\beta_j - \hat{\beta}_j^{(i)})\right]\right\} \\
&= -\frac{(N+1)m}{2}\log(2\pi) - \frac{N+1}{2}E_{q_i}(\log|\Sigma|) \\
&\quad - \frac{m}{2}\log|(V_j^{(i)})^{-1}| - \frac{1}{2}m(N+1)
\end{aligned}
$$

$$
E_{q_i}\left\{\log p(y_i|\theta_{\delta_i})\right\} = \sum_{j=1}^{T}\hat{q}_i(\delta_i = j)E_{q_i}\left\{\log(y_i|\theta_j)\right\}
$$

where

$$
\begin{aligned}
E_{q_i}\left\{\log p(y_i|\theta_j)\right\} &= -\frac{m}{2}\log(2\pi) - \frac{m}{2}E_{q_i}\left\{\log\tau\right\} - \frac{1}{2}E_{q_i}\left\{\log|\Sigma|\right\} \\
&\quad - \frac{1}{2}E_{q_i}\left\{\frac{1}{\tau}\right\}E_{q_i}\left\{y_i^T\Sigma^{-1}y_i - 2E_i^T\beta_j\Sigma^{-1}y_i + E_i^T\beta_j\Sigma^{-1}\beta_j^T E_i\right\} \\
&= -\frac{m}{2}\log(2\pi) - \frac{m}{2}\left\{\log(b_\tau^{(i)}) - \psi(a_\tau^{(i)})\right\} \\
&\quad - \frac{1}{2}\left\{-\psi_m\left(\frac{\nu^{(i)}}{2}\right) - m\log(2) + \log|S^{(i)}|\right\} \\
&\quad - \frac{1}{2}\frac{a_\tau^{(i)}}{b_\tau^{(i)}}\left\{\nu^{(i)}y_i^T(S^{(i)})^{-1}y_i - 2\nu^{(i)}E_i^T\hat{\beta}_j^{(i)}(S^{(i)})^{-1}y_i\right. \\
&\quad \left. + mE_i^T\Omega_j^{(i)}E_i + (E_i^T\hat{\beta}_j^{(i)})^T(\nu^{(i)}(S^{(i)})^{-1})(E_i^T\hat{\beta}_j^{(i)})\right\}.
\end{aligned}
$$

Finally, we have

$$
E_{q_i}\left\{\log(q_i(\delta_i)\right\} = \sum_{j=1}^{T}\hat{q}_i(\delta_i = j)\log(\hat{q}_i(\delta_i = j)).
$$

References

Attias, H. (2000). A variational Bayesian framework for graphical models. In *Advances in Neural Information Processing Systems 12*, pp. 209–215. MIT Press.

Beaumont, M. A., W. Zhang, and D. J. Balding (2002). Approximate Bayesian computation in population genetics. *Genetics 162*, 2025–2035.

Blei, D. M. and P. I. Frazier (2011). Distance dependent Chinese restaurant processes. *J. Mach. Learn. Res. 12*, 2461–2488.

Blei, D. M. and M. I. Jordan (2006). Variational inference for Dirichlet process mixtures. *Bayesian Anal. 1*, 121–143.

Blum, M. G. and V. C. Tran (2010). HIV with contact tracing: A case study in approximate Bayesian computation. *Biostatistics 11*(4), 644–660.

Blum, M. G. B. (2010). Approximate Bayesian computation: A nonparametric perspective. *Journal of the American Statistical Association 105*(491), 1178–1187.

Blum, M. G. B. and O. François (2010). Non-linear regression models for approximate Bayesian computation. *Statistics and Computing 20*, 63–75.

Box, G. (1980). Sampling and Bayes' inference in scientific modelling and robustness (with discussion). *Journal of the Royal Statistical Society, Series A 143*, 383–430.

Bryant, M. and E. B. Sudderth (2012). Truly nonparametric online variational inference for hierarchical Dirichlet processes. In *Advances in Neural Information Processing Systems 25*, pp. 2708–2716.

Caron, F., M. Davy, and A. Doucet (2007). Generalized Pólya urn for time-varying Dirichlet process mixtures. In R. Parr and L. C. van der Gaag (Eds.), *UAI*, pp. 33–40. AUAI Press.

Chen, C., V. Rao, W. L. Buntine, and Y. W. Teh (2013). Dependent normalized random measures. In *ICML (3)*, Volume 28 of *JMLR Proceedings*, pp. 969–977.

De Iorio, M., P. Müller, G. L. Rosner, and S. N. MacEachern (2004). An ANOVA model for dependent random measures. *Journal of the American Statistical Association 99*, 205–215.

Dunson, D. B. and J.-H. Park (2008). Kernel stick-breaking processes. *Biometrika 95*, 307–323.

Evans, M. and G. H. Jang (2011). Weak informativity and the information in one prior relative to another. *Statist. Sci. 26*(3), 423–439.

Evans, M. and H. Moshonov (2006). Checking for prior-data conflict. *Bayesian Anal. 1*(4), 893–914.

Foti, N. and S. Williamson (2015). A survey of non-exchangeable priors for Bayesian nonparametric models. *IEEE Transactions on Pattern Analysis and Machine Intelligence 37*, 359–371.

Gelfand, A. E., A. Kottas, and S. N. MacEachern (2005). Bayesian nonparametric spatial modeling with Dirichlet process mixing. *Journal of the American Statistical Association 100*, 1021–1035.

Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models. *Bayesian Anal. 1*(3), 515–533.

Gelman, A., A. Jakulin, M. G. Pittau, and Y.-S. Su (2008). A weakly informative default prior distribution for logistic and other regression models. *Ann. Appl. Stat. 2*(4), 1360–1383.

Griffin, J. E. and M. F. J. Steel (2006). Order-based dependent Dirichlet processes. *Journal of the American Statistical Association 101*(473), 179–194.

Hoffman, M. D., D. M. Blei, C. Wang, and J. Paisley (2013). Stochastic variational inference. *J. Mach. Learn. Res. 14*, 1303–1347.

Jordan, M. I., Z. Ghahramani, T. S. Jaakkola, and L. K. Saul (1999). An introduction to variational methods for graphical models. *Mach. Learn. 37*, 183–233.

Kabisa, S. T., D. B. Dunson, and J. S. Morris (2016). Online variational bayes inference for high-dimensional correlated data. *Journal of Computational and Graphical Statistics To appear*.

Kingman, J. F. C. (1967). Completely random measures. *Pacific Journal of Mathematics 21*, 59–78.

Lijoi, A., B. Nipoti, and I. Prünster (2014). Bayesian inference with dependent normalized completely random measures. *Bernoulli 20*, 1260–1291.

Lijoi, A. and I. Prünster (2010). Models beyond the Dirichlet process. In N. L. Hjort, C. Holmes, P. Müller, and S. G. Walker (Eds.), *Bayesian Nonparametrics*, pp. 80–136. Cambridge University Press.

Lin, D. (2013). Online learning of nonparametric mixture models via sequential variational approximation. In C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger (Eds.), *Advances in Neural Information Processing Systems 26*, pp. 395–403. Curran Associates, Inc.

Luts, J., T. Broderick, and M. P. Wand (2014). Real-time semiparametric regression. *Journal of Computational and Graphical Statistics 23*, 589–615.

MacEachern, S. N. (2000). Dependent Dirichlet processes. Technical report, Department of Statistics, Ohio State University.

Nott, D., C. C. Drovandi, K. Mengersen, and M. Evans (2015). Approximation of bayesian predictive p-values with regression abc. Technical report, Queensland University of Technology.

Ormerod, J. and M. Wand (2010). Explaining variational approximations. *The American Statistician 64*, 140–153.

Racine, A., A. P. Grieve, H. Flühler, and A. F. M. Smith (1986). Bayesian methods in practice: Experiences in the pharmaceutical industry. *J. Roy. Statist. Soc. Ser. C 35*(2), 93–150.

Rao, V. A. and Y. W. Teh (2009). Spatial normalized gamma processes. In Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, and A. Culotta (Eds.), *Advances in Neural Information Processing Systems 22*, pp. 1554–1562.

Sato, M.-A. (2001). Online model selection based on the variational bayes. *Neural Comput. 13*, 1649–1681.

Teh, Y. W., M. I. Jordan, M. J. Beal, and D. M. Blei (2006). Hierarchical Dirichlet processes. *Journal of the American Statistical Association 101*(476), 1566–1581.

Tsanas, A. and A. Xifara (2012). Accurate quantitative estimation of energy performance of residential buildings using statistical machine learning tools. *Energy and Buildings 49*, 560 – 567.

Wang, C. and D. M. Blei (2012). Truncation-free online variational inference for Bayesian nonparametric models. In F. Pereira, C. Burges, L. Bottou, and K. Weinberger (Eds.), *Advances in Neural Information Processing Systems 25*, pp. 413–421. Curran Associates, Inc.

Wang, C., J. Paisley, and D. M. Blei (2011). Online variational inference for the hierarchical Dirichlet process. In *Proc. of the 14th Int'l. Conf. on Artificial Intelligence and Statistics (AISTATS)*, Volume 15, pp. 752–760.

Wang, L. and D. B. Dunson (2011). Fast Bayesian inference in Dirichlet process mixture models. *Journal of Computational and Graphical Statistics 20*, 196–216.

Waterhouse, S., D. Mackay, and T. Robinson (1996). Bayesian methods for mixture of experts. In *Advances in Neural Information Processing Systems 8*, pp. 351–357. MIT Press.

Zhang, X., D. J. Nott, C. Yau, and A. Jasra (2014). A sequential algorithm for fast fitting of Dirichlet process mixture models. *Journal of Computational and Graphical Statistics 23*, 1143–1162.

Zhang, Z., G. Dai, and M. Jordan (2010). Matrix-variate Dirichlet process mixture models. *Proceedings of the Thirteenth Conference on Artificial Intelligence and Statistics (AISTATS) 9*, 988–995.

Zhang, Z., D. Wang, G. Dai, and M. I. Jordan (2014). Matrix-variate Dirichlet process priors with applications. *Bayesian Analysis 9*, 259–286.